Bob Hanké

How including labour can improve corporate governance

Book section (accepted version)

Original citation:
Originally published in Driver, C.; Thompson, G., *Corporate governance in contention*. Oxford: Oxford University Press, 2018, 226-238.

(in) Driver, Ciaran and Thompson, Grahame (eds.), *Corporate Governance in Contention,* Oxford University Press, 2018: 226-38

How including labour can improve corporate governance

Bob Hancké, LSE

June 2017

Abstract

Involving labour in decision-making has the potential to improve corporate governance, even in adversarial industrial relations systems such as the ones found in Anglo-American economies. Think of corporate governance as an information problem with potentially dramatic distributive consequences: how do shareholders and other interested parties to the activities of a company know that management is working in their best long-term interest? Without labour on board as a vocal actor in corporate decision-making, profits, often defined as short-term gains, will trump other considerations, while a system in which labour holds veto powers is likely to lead to lower profits, ceteris paribus. If both labour and business are represented in decision-making, however, the information asymmetries that each faces are significantly alleviated by the presence of the other, which leads to more balanced outcomes. Representatives of business – financial institutions, suppliers, or other firms in a similar industry know relatively little about how a company is run, but a lot about how the company is doing in its key product markets. Labour, on the other hand, may have only a tenuous grip on competitive strategy, but is highly cognisant of how the company is run – it has to deal , after all, with the problems that arise. A board system, not unlike the north-west European one, where representation is shared among these two key actors, thus forces management to be transparent and take into account the preferences of both.

Acknowledgments

Comments from Ciaran Driver and Grahame Thompson are gratefully acknowledged, as are broader discussions over the years with Richard Bronk, Saul Estrin, Michel Goyer and David Soskice over the years. The usual exculpations apply. Comments welcome: r.hancke@lse.ac.uk

What can workers and their representatives (henceforth labour) bring to corporate governance and, through that channel, contribute to company performance (Estrin & Marsden 2017)? For some this may be a somewhat awkward question: labour is, after all, considered to be a countervailing power to capital, and in the conventional view representation in strategic decision-making bodies in the company exists to safeguard labour's rights, further labour's agenda and generally reduce unilateral managerial power. But even a cursory glance at north-western Europe, where the institutionalized presence of labor in company decision-making is widespread – and where wages are high, skills well-developed, and companies easily find new places in existing industries while social conflict is near absent – suggests that the link between labour and corporate decision-making is more complex.

In this chapter, I will make a simple analytical point about the benefits for corporate governance of having representatives of both shareholders and employees (also called capital and labour henceforth) in strategic decision-making – either on the Board of Directors, as in Anglo-American systems, or on the supervisory board, as in continental systems. My argument is that the simultaneous presence of both sides of industry results in two very different but complementary information sets being combined: one on the 'external' situation of the company in its markets and vis-à-vis its competitors, brought in by capital; and another, on the internal operations of the company through labour. As a result, managers are permanently caught in a situation where at least one party on the board has at least as much information, and usually more, than they do, which makes it impossible to exploit the informational advantage that management would otherwise have.

I will start with a quick review of the key model in corporate governance – the principal-agent model – how it conceptualizes the problem, and what is wrong with the solutions it proposes. The paper then goes on to examine an employee-inclusive corporate governance system along the lines of what exists in some form or other in many continental north-west European countries and how this reduces informational asymmetries. This is followed by a short case study of Co-

Determination in Germany and its business effects. I conclude with a few thoughts on wider conditions that have to be met for such systems to be successful.

Before proceeding, two caveats are in order. One, I will concentrate on the basic information asymmetry problems that haunt modern corporate governance, and much less on the problem of cooperation between board members and management (see Adams & Ferreira 2007). There is some scattered evidence that most non-executive directors are on (too?) many boards and in fact not very knowledgeable about the company or the industry they are working in; even though I treat them as competent directors below, it is clear that their contribution to management might be somewhat limited as a result. Secondly, I will focus on the instrumental value of such a redesigned corporate governance system, and only concentrate on the bottom line of a company, without addressing the benefits to wider stakeholders, positive reputational effects of good behaviour, or normative arguments about the nature of the capitalist firm; the model I propose is perfectly compatible with such ideas, but would unnecessarily broaden the discussion for my purposes.

Corporate governance: principals and their agents

The standard approach to the problem of corporate governance is known as the principal-agent model: owners and their representatives (the principal) hire a professional manager (an agent) to run the company in their best interest, i.e. in such a way that the value of the company increases, expressed in rising share values, dividends or both. In principle owners are agnostic about how this is achieved, though increasingly secondary considerations have imposed relatively harder constraints (viz. corporate social responsibility). The problem with this set-up is that the manager has access to information that the owner has not and might be able to use this information to divert resources from the core objective of increasing the value of the company. Put differently, the principal knows less, but does not know what exactly is unknown and is therefore unable to adequately monitor and evaluate the actions of its agent and their potential implications.

The standard solution to this problem is to align the incentives of the agent and the principal: a significant part of the CEO's remuneration is tied to those elements in the performance of the company that also benefit the owners. The classic example is a deferred salary in the form of stock options: since both the owner's and the manager's income increases if the share price rises, owners are assured that management efforts also reflect their interest. Another version is one in which management is set hard targets, leading to a salary bonus when met; or the bonus is granted afterwards on the basis of individual or group performance. All these mechanisms have in common that management is disciplined because of the public nature of the information: share prices, targets, and performance indicators, regardless of the nature of the private information that management holds, can relatively easily be gathered and evaluated.

But this arrangement is not without problems (Bronk 2009). First of all, as we saw in the run-up to the financial crisis, the focus on shareholder value and targets, rewarded by performance bonuses, created a series of adverse (and often perverse) incentives[1]. Mortgage sellers were, as we discovered, rewarded on the basis of the number of mortgages they generated, rather than the 'quality' of the loan. Borrowers, on the other hand, were encouraged to take out mortgages well beyond what their income could afford, in the knowledge (and hope) that the rise in value of the property would offset the gap between what was affordable and the mortgage repayments. Rating agencies, which also served the mortgage banks, faced strong incentives to award their customers and their products higher ratings than they might have deserved. That implied that for other banks, such highly rated bonds were considered more or less risk-free, which encouraged the construction of a market. Performance bonuses, in turn, were notoriously short-term, and usually only covered the past year, regardless of where the mortgages and the bonds went afterwards (and we all know how that actually ended). In short, what appeared to be a rational system of incentive alignment turned out to be a snake pit riddled with information asymmetries and worse – but unchecked by the system as a whole.

It is not that difficult to see how these problems, which took their almost paradigmatic form as causes of the financial crisis of the 2000s, also exist in some form beyond the financial sector. If management's salary is disproportionately tied

[1] My colleague Nick Barr set out some of these in a talk to our students at the LSE a few years ago and I am very grateful to him for allowing me to borrow them for my purpose here.

to relatively short-term performance indicators, the incentives are such that longer-term goals, which might be more lucrative for the company (but, crucially, not necessarily for the individual manager), will be sacrificed to possibly detrimental short-term objectives that raise the share price.

Secondly, the winner-takes-all labour market of top management has effectively split 'management' into a very small group – in the limiting case only the CEO – where high-powered incentives are focused, and a larger group of high-end employees, who might have detailed knowledge of what is going on in the company but are not directly involved in running the company and reporting to shareholders. That is an important source of experience and relevant information that is increasingly lost in contemporary business.

The third problem with such incentive schemes is that they do not, or at best only imperfectly, take into account wider effects on other groups that are affected by them. Whilst such negative externalities are deemed not to exist in the theory of perfect markets at the basis of these corporate governance models, in practice they are often significant and – most importantly – not symmetrically distributed across potentially affected groups. Take a simple example: if company results are flagging, management often will take recourse to cost-cutting measures – investment is often the first point of call, usually followed by actions to reduce the total wage sum. But the first leads to a cascade down the line: the machine tool builders may not survive a downturn (and, if these were highly specialized and dedicated machine tool builders, their demise may deprive the company of its next generation of machine tools). The wage cost reduction leads either to a fall in consumption, unemployment among the workforce and, in the longer run, to an erosion of the implicit social contract between a company and its workforce to invest in skills that are realized in the long term. Often it is not even possible to calculate the negative externalities in any meaningful detail: when VW introduced its 28-hour week in the mid-1990s in response to a profitability crisis, the social and economic life of the Wolfsburg area – from football clubs to kindergarten and restaurants – took a massive shock.

Fourth and finally, aligning incentives between managers and owners along the lines above leads, in practice, often to an overemphasis on short-term thinking. While this is certainly not a logical necessity, the fact that it often appears suggests that something deeper is at work here. In part, this bias in favour of the short term must be related to the fact that there is simply no satisfactory analytical definition of the long term, as Keynes implied in his famous quip. But in part it also follows from the incentive structures themselves: nothing improves future job chances for a manager in another company as much as the demonstration of the ability to raise the share price now. Indeed, this suggests that the short-termism is deeply embedded in the architecture of corporate finance. In effect, shareholders reward a CEO's short-term orientation because of their almost exclusive attention to the share price. And, while this is particularly the case in the Anglo-American world, the search for capital on international markets has spread these job-hopping practices for top executives, even in the once more relaxed systems in northern continental Europe.

This leaves us in a quandary. If the standard solution of aligning the incentives of owners and managers does not resolve the issues associated with deep information asymmetries, and if this produces a host of additional problems, how can we resolve this? In the section below, I want to introduce a mechanism that avoids or at the very least mitigates these tensions. The basic idea is that introducing two symmetrically opposite but complementary information sets in the corporate governance structure of companies, one associated with capital and another associated with labour, monitoring and disciplining management is possible without introducing perverse effects.

Capital and labour in corporate governance: complementary information sets

Let us start this exploration by thinking of corporate governance not just as a narrow matter of owners and managers, but more widely as the mechanism that assures decisions are made which balance different claims, ownership-related or otherwise, on the company. Ownership remains important in this perspective, but is embedded in a wider web of relations between the company and its environment: companies make profit, but they also produce other outcomes, desired or not, for which the company often bears little liability – externalities which result from an incomplete assignment of property rights. Assume also, for the sake of the argument, that we are looking at the best possible world, in which both CEOs and company directors are competent and put in effort in their management and overseer roles. For the purposes of this analysis, I will limit the wider context to the two key factors of production in modern industry: capital and labour. Decision-making in companies follows a system in which capital and labour both have between one third and half of the seats on the board, and the person chairing the board, usually from the capital side, has a casting vote if necessary. This is roughly speaking the situation that exists in countries like Germany, Austria, Sweden and other northern continental European economies. The set-up of the decision-making structure is such that adversarial votes are avoided (not least since trade unions have hard rights outside these forums which could upset any unilateral management initiatives), and that both sides prefer to go the slow route by bargaining out a mutually viable solution. Decisions may be slower to reach because of the implied obligation to negotiate, but also easier to implement afterwards since all relevant actors are on board.

It might be important to emphasize that in such a system, attention to profits does not disappear; it is also in the interest of the workforce that a company does well. What might and almost certainly will change is the distribution of the surplus between returns to owners, to workers as wages, and investment in skills and technology, depending on the nature of the bargain between the different parties. Management can get on with managing, but simply has to take more constraints into account. But what will almost certainly disappear is the often unhealthy obsession with short-term results. Since costs as well as gains for different actors are actively incorporated in every decision, they will balance the different time horizons for each of the parties. All other things equal, therefore, a joint capital and labour decision-making structure will be geared towards long-term, often more sustainable, solutions; in fact, many problems will ordinarily not arise or develop, since they will have been picked up before they reach a critical stage.

It is probably useful here to draw a distinction between companies in industries that require relatively quick responses to turbulent markets or rapid technological shifts, on the one hand, and those in more mature industries with relatively stable technologies on the other. In the former, a 'thick' decision-making model that involves labour is probably not as good an idea as in the latter, since it unnecessarily limits the CEO's ability to adjust rapidly, while the potential losses for employees and for the company are relatively small, since the knowledge of the workforce is intrinsically part of strategic decision-making in most of these highly dynamic industries anyway. It is hard to imagine a cutting-edge biotechnology company or a software developer without a strategic role for its engineers. But markets and technologies eventually stabilize and reach a point where management can more easily plan rather than primarily respond to outside shocks (Casper 2007: xxx-xxx). Once new molecules exist, the key problem is to stabilize and develop them into commercial products, a considerably more 'bureaucratic' process than discovering or synthesizing them. And there is nothing that precludes introducing employee participation in decision-making once that stage in the product life cycle is reached.

Negative externalities, especially those that directly affect the workforce, are minimized when labour is an active partner in decision-making. In the longer-term decision-making context that characterizes this set-up, which could be understood loosely as an iterative cooperative game, losses today are likely to be compensated by gains tomorrow. It is highly unlikely that one of the two parties yields a permanent structural advantage over the other. This is reinforced by the nature of labour markets and product markets where such companies are positioned: the long-term horizons allow them to seek out relatively safe product market niches, fed by deep occupation-based skills, which gain in value for these product market strategies as they are developed. This corporate governance arrangement, which includes labour and capital, thus avoids several of the drawbacks associated with the standard principal-agent model.

Most importantly, however, a corporate decision-making structure that includes capital and labour also circumvents the information asymmetries that are endemic in the standard principal-agent model. Starting from the existing, one-sided information context that we know well from Anglo-American economies as a baseline scenario helps understand how it does so. An external control model, which runs through ownership of company shares, faces two types of related information problems. One is that managers have a structural advantage related to its monopoly over strategic information; in effect, as Roe (1994) argued over two decades ago, the outcome combines strong managers with weak owners – almost the opposite of what it is supposed to do under standard corporate governance theory. Now, since no single owner has a strong incentive to closely monitor management because individual stakes in the company are too small, they usually rely on a series of relatively transparent, public performance indicators; and instead of exercising voice, owners sell shares when they deem company performance below par. The only disciplining mechanism left is the market for corporate control: if the share price falls too sharply, a hostile takeover will not only eliminate the company as a separate entity, it is almost certain to entail the dismissal of the underperforming manager (often involving, somewhat ironically, a substantial redundancy package for the failing individual). Whilst the possibility of a takeover will certainly focus the minds of managers, such a nuclear option is effectively rarely exercised: if the mechanism matters, it does so primarily by acting as a warning to management about the importance of performance, not as an ex post sanctioning mechanism.

With both sides at the table, however, things change significantly. The informational advantage that management has over each is reduced dramatically or disappears entirely as a result – and this is the important bit – of the information offered *by the opposite side*. Imagine a meeting of the board, where many of the shareholders and directors are aware of problems with the company: they see how its main competitors are doing and discover that its results, expressed in profits, share price or market share, are weak and show no signs of improving. Holding management to account sounds easy but may not be, since the manager can claim that this is a temporary problem resulting from an internal reorganization, that reinvestment plans are pending, that reducing costs is difficult, or that this is likely to happen before a major restructuring of operations. Board members have very little inside information that they are able to mobilize to evaluate these claims. Conversely, in negotiations with the workforce, claims by management that the company will be doing much better in two years, but that higher wages and increased training are impossible now, because the main competitors have launched a campaign to increase market share.

As long as management can talk to each individually, it is very difficult for any of the two parties, shareholders and workforce, to counter these claims without a significant investment in gathering information which is not very easily accessible (else the information problem would, by definition, not exist). Together, however, they will often master a large part of the relevant information and, most importantly, in a way that compensates for the holes in the information set of the other side. Directors, shareholders and other board members usually have a very good sense of, for example, how markets are structured, how technology generally is developing, how well the company is doing compared to others in the industry, or if the industry in which the company is active has a vibrant future. Put differently, representatives of owners are, by virtue of their position in industry and related activities, very knowledgeable about the external position of the company and its performance. At the same time, however, they know relatively little about how the company is faring internally: whether management is facing significant internal tensions, if skill development is keeping pace with other companies, if wages are commensurate with skills and reflect industry norms, and if turnover is high, for example. Labour, in turn, faces exactly the opposite problem (with, perhaps, the exception of a small group of high-level middle managers): it has virtually no knowledge about markets, technology, competition and the like but is very well aware of how the company is operating internally – if only because labour's representatives are usually the ones who respond to problems that might exist with that.

Corporate governance systems in which representatives of both capital and labour are board members in effect combine two highly developed, complementary information sets that take management in a pincher movement. Issue areas where capital does not have sophisticated knowledge are likely to be among those on which labour has superior (inside) knowledge; fields in the performance of the company where labour has at best a tenuous grasp of the problem, are likely to be among those where capital has a significant informational advantage. Together, they cover those areas of company organization and performance that allow them to evaluate management performance on the basis of deep and relevant information. The informational advantage that management has as a result of its privileged position has all but disappeared.

German Co-Determination

The existing model closest to this abstract, almost ideal-typical corporate governance arrangement that includes capital and labour in strategic decisions is found in large firms in Germany. It involves a dual board structure, with a management committee (of more or less equals) consisting of the CEO, the COO and the 'Labour Director', and a supervisory board (*Aufsichtsrat*), which offers strategic guidance and monitors management performance. The supervisory board comprises representatives of management and labour, in a representation system with staggered thresholds according to company size. In companies from 500 employees up, minimum one-third of the supervisory board seats are reserved for labour (including one for middle management), while in companies with a workforce above 2000, the seats are distributed 50-50, usually chaired by a representative of the shareholders, with a casting vote in case of a tie. Of the workforce representatives, about half are from the company workforce, while the other half are appointed by the representative union in the industry (Germany follows, in most cases the 'one industry – one union' principle). All parties have access to the same internal company information and can draw on a wide array of outside experts to help them in their task of company director.

Surprisingly, perhaps, especially from an Anglo-American perspective, evaluations by outsiders, including academics, have been quite positive. In the late 1960s and early 70s, and again in the 2000s, two Commissions chaired by the same Christian-Democratic politician Kurt Biedenkopf concluded that this system of corporate governance, known as Co-Determination (*Mitbestimmung*), at the very least did not produce any significant costs for management while offering benefits such as social peace and cooperative workplace industrial relations. In fact, despite occasional dissident noises from employers associations, including in the latest Biedenkopf Commission, a majority of managers are in favour of retaining the existing institutions (Silvia 2013).

The history of Co-Determination shows that this is far from an uncontested system (Jackson 2005), and that it therefore regularly needs to prove itself not just on the terms in which it was initially conceived – giving workers a meaningful say in how industry is run – but also in terms of company results (Silvia 2013: 43-64 offers an excellent review of this history). In a potted version, the post-war history started in the 1950s with a general law, followed by revisions in the 1970s and 1980s, when the 2000-employee threshold for a 50% representation was introduced, and some minor changes in the 2000s. On each of these occasions, employers sought to curtail employees' influence in decision-making on the predictable grounds that it hindered management autonomy, while unions fought hard to keep the arrangements roughly as they were based on equally predictable arguments of social justice. That said, after each of the rounds of changes (the early ones in the 1950s went rather significantly against the unions, who had asked for more far-reaching employee representation), all parties accepted the arrangement and worked within the new rules.

Most of the assessments of the economic effects of Co-Determination are quite positive. A first indicator is that the risk premium on corporate bonds for large firms headquartered in Germany, and therefore subject to Co-Determination rules, and outside Germany, where this is not the case, is insignificant (HBS 2007). The same applies to FDI, where Germany does not underperform compared to other G7 economies (HBS 2007). Secondly, the performance of companies in Germany employing more than 2000 employees, where employees hold 50% of the SB seats, is roughly similar to those with a workforce below 2000, where representation is limited to one third (Vitols 2006: 12). The differences in return on investment between the former and the latter are small and variations over time go in both directions. Similarly the ratio of share price to book value is, in fact, higher in the larger companies than in the smaller ones. Finally, there are strong indications that job satisfaction and productivity are higher in companies with more than 2000 employees than in those with a workforce between 1500 and 2000 (Loose et al. 2011). Taken together, these data suggest quite strongly that the potential management penalty associated with Co-Determination is non-existent or, at worse, very small, and that the best performing companies in Germany are those that see Co-Determination not as a source of costs and other problems, but as a strategic asset to be deployed for the benefit of the company.

Background conditions

Despite their potential benefits, such systems of joint governance are difficult to build and possibly equally hard to sustain. It is also unclear to what extent the absence of workers' voices in corporate governance is, as it were, an accident of history that can be corrected, and to what extent different corporate governance systems are simply incompatible with effective institutionalized employee voice. I have adopted the moderately optimistic position here that such institutional transplants are, in principle, possible – but that then raises the question of conditions under which it could develop. This final section reviews a small number of what appear to be necessary conditions to make such a system work.

One of the key conditions for such a joint decision-making system to function well is a measure of mutual trust, in which both sides accept to take risks in the understanding that the other party will not exploit any vulnerabilities that emanate from committing to cooperation. Since such trust grows over many years of interaction – not necessarily from a consensus, but usually from struggles over stewardship of the economy and companies, which is then turned into a strategic asset – building trust without such historical foundations can be very difficult.

Adversarial industrial relations systems, such as the ones found in Anglo-American economies and in 'ideological' labour relations systems (as in southern Europe, with politically organized trade union confederations) are probably not particularly well positioned in this respect. They may have to find a way around too many obstacles as a result of their history. Yet building trust is not necessarily impossible, however, even from such a contrarian starting point. It might be useful here to conceptualize trust a bit more carefully, and distinguish between blind trust and 'studied' trust (Sabel 1993): the first is a situation in which A has no other option than to leave their fate in the hands of B because they are extremely dependent (the paradigmatic case is that of a toddler and a parent) or alternatives are simply unavailable; the second form of trust is best understood as built on positive feedbacks that result from previous cooperation. The development of local and regional training systems is an example of the latter, in which the trust engendered by positive effects of small experiments becomes the basis for a broader discussion among business, employer associations, local chambers of commerce and government, and employees and their representatives. Since trust grows in value as it is used, any positive experiences are likely to produce further (and probably more intensive) subsequent cooperation. While such discussions in narrow areas will start out limited to issues where the key parties see immediate joint benefits, other problem areas are likely to become part of a widening agenda as the system evolves. Studied trust captures a situation in which parties discover, even without long histories of mutual adjustment and cooperation, the benefits of rewarding cooperative behaviour – which then unleashes an endogenous cooperative dynamic (see Hancké 2012 for an illustration of this in Central Europe, a region with 'institutionally thin' political economies not unlike the UK).

Deep cooperation of this kind requires competent actors on all sides of the table. The northern continental European systems that were the explicit or implicit benchmark for the analysis in this chapter have large budgets for training labour representatives and independent expertise. In Germany, for example, trade union members of the supervisory boards donate the fees they are paid for their work on the supervisory board to the trade union confederation, which then invests it in research and training programmes. At the very least, a system in which workers perform roles on boards would require a similar organizational and financial backup.

Yet such expertise can turn meaningless if it is centrally provided. Take an example from another area in political economy: the Monetary Policy Committee of the Bank of England, where every member has access to their own (small) research team, and can thus provide an independent assessment of the future evolution of inflation. Since not every member of the MPC works with the same underlying model of the economy, they will come up with different views. Combined, however, these independent individual assessments will provide a clear view of possible scenarios and their implications – and thus improve decision-making. Transposed into the area of corporate governance, the independence of the expert advice (and training provisions) is equally important: if the company gains from having the views of workers' representatives on the board which are different from those of owners, and if workers' representatives indeed concentrate their efforts in different areas than management and owners as I suggested earlier, it makes sense to provide them with their own independent training and expertise capabilities – obviously within a broader framework of loyalty to the company. It is considered standard practice in

German large firms, for example, that unions and workers representatives develop and submit alternative plans in the case of a large-scale restructuring that will work towards managing costs within the same envelope as management plans, but where the employment effects are mitigated. Without independent expert advice such constructive alternatives are probably impossible to develop.

Conclusion

If owners face a structural information asymmetry vis-à-vis managers, and if a simple alignment of the incentives of managers and owners does not resolve these information problems, a solution might reside in bringing in workers' representatives on boards. The basic idea behind this chapter is simply that representatives of owners and workers offer different but complementary perspectives on the management of the firm. Whereas board members that represent owners will have a relatively good sense of company strategies in key markets but a weak sense of internal organization, workers' representatives face the opposite problem of deep knowledge about the internal functioning of the company but a much weaker sense of how the company fares in its market environment. A set-up where employees are also represented will not only allow for a more pointed monitoring of management, but will also provide management with relevant helpful information in the case of problems. Having set out this argument, this chapter presented a short case study of German co-Determination, one of the most effective and institutionalised forms of joint decision-making in the world, and evaluated its effects in a handful of essential dimensions of company performance. I finished the analysis by reflecting on a series of conditions that have to be met for such information flows from workers' representatives to perform this function – of which independent expertise and training seem paramount.

We know from experience in continental Europe that a joint or parity-based representation system can work well. It is hard to make the case that actively listening to views from different sides of industry has handcuffed management in German, Swedish, or Austrian companies and potentially jeopardised their performance. It is not even clear if in the long run any costs associated with having workers' representatives on boards are not far outweighed by their benefits. Secure in the knowledge that their interests are protected in strategic decisions, workers are likely to invest more in training and deploy their knowledge for the good of the company, engage constructively with restructuring plans, and refrain from strikes. Decision-making may be slower, but solutions effectively gain in legitimacy and are therefore easier to implement.

But institutions do not just emerge when they solve problems. They might eventually become part of the solution, but they usually emerge under conditions of institutional reconfiguration, often as a result of struggles. As the section on the German system suggested, none of the northern European models in existence reflected the preferences of business when introduced, and it is important to bear this contingent nature of the institutional innovation in mind (Fetzer 2010; Jackson 2005). Yet it is equally clear that without shared benefits any new decision-making structure in companies is likely to fail in attracting the support it needs. In this chapter I have offered a way of making sense of these benefits of joint decision-making arrangements. With that hurdle cleared, political mobilisation can put it on the agenda. Employees deserve a stronger voice and companies will be better run as a result.

Bibliography

Adams, R.B. & Ferreira, D. 2007. 'A Theory of Friendly Boards', *The Journal of Finance* (62:1), 217-50.

Bronk, R. 2009. *The Romantic Economist : Imagination in economics*. Cambridge, UK: Cambridge University Press.

Casper, S. 2007. *Creating Silicon Valley in Europe : Public policy towards new technology industries*. Oxford: Oxford University Press.

Estrin, S. & Marsden, D. 2017. LSE Department of Management Response to DBEIS Green Paper on Corporate Governance Reform, Manuscript, Dept of Management, London School of Economics and Political Science.

Fetzer, T. 2010 'Defending Mitbestimmung: German Trade Unions and European Company Law Regulation (1967-2000).' *Economic and Industrial Democracy* 31.4_suppl: 24-39.

Hancké, B. 2012. Endogenous coordination: Multinational companies and the production of collective goods in central and Eastern Europe. Stato e Mercato: Quadrimestrale di Analisi Dei Meccanismi e Delle Istituzioni Sociali, Politiche ed economiche , 2012 (2). pp. 203-227.

HBS. 2007. ‘Results of the “Biedenkopf Commission”: Government Commission on the modernization of employee board-level representation in Germany. Executive Summary’, Duesseldorf: Hans-Boeckler-Stiftung (https://www.boeckler.de/pdf/bb_zusammenfassung_BiKo_engl.pdf; accessed on 16 June 2017)

Jackson, Gregory. ‘Employee Representation in the Board Compared: A Fuzzy Sets Analysis of Corporate Governance, Unionism and Political Institutions.’ *Industrielle Beziehungen / The German Journal of Industrial Relations* 12.3 (2005): 252-79.

Looise, J.K., Torka, N. & Wigboldus J.E. 2011. ‘Understanding Worker Participation and Organizational Performance at the Firm Level: In Search of and Integrated Model’, *Advances in Industrial and Labor Relations*, 18,87-113.

Roe, M. (1994). *Strong managers, weak owners : The political roots of American corporate finance*. Princeton, N.J. : Princeton University Press 1994

Sabel, Charles. "Studied Trust: Building New Forms of Cooperation in a Volatile Economy." *Human Relations* 46.9 (1993): 1133

Silvia, Stephen J. 2013. *Holding the Shop Together: German Industrial Relations in the Postwar Era*. Ithaca NY: Cornell University Press.

Vitols, S. 2006. *Oekonomische Auswirkungen der paritaetischen Mitbestimmung: Eine oekonomische Analyse,* Duesseldorf: DGB