IZA Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 10671

Past, Present and Future Compensation Research: Economist Perspectives

Michael Gibbs

MARCH 2017

IZA Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 10671

Past, Present and Future Compensation Research: Economist Perspectives

Michael Gibbs
University of Chicago and IZA

MARCH 2017

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Past, Present and Future Compensation Research: Economist Perspectives*

At the 2016 Academy of Management Conference, a group of distinguished compensation researchers held a panel discussion on the future of compensation research. Their remarks were compiled into an article published in this issue. Soon after the panel, Charles Fay commissioned a similar discussion by leading economists (Michael Gibbs, Kevin Hallock, Edward Lazear, Kevin J. Murphy & Canice Prendergast) doing compensation research; this article is the result.

JEL Classification: J3, M5

Keywords: compensation, incentives, evaluation, intrinsic motivation

Corresponding author:
Michael Gibbs
University of Chicago Booth School of Business
5807 S. Woodlawn
Chicago, IL 60637
USA

E-mail: michael.gibbs@chicagobooth.edu

* Special thanks to Kevin Hallock, Ed Lazear, Kevin J. Murphy and Canice Prendergast, who generously provided their thoughts and valuable time. I also thank Charles Fay for suggesting this panel discussion. Funding from the University of Chicago Booth School of Business is gratefully acknowledged.

Labor economists have studied compensation in great detail for a very long time. Most of that research was from the perspective of labor markets, or of individual workers. For example, extensive literatures exist on how pay varies with employee skills, education, and demographic characteristics. Labor economists have studied how wages vary between private and public sectors, and between union and non-union firms. Macroeconomists have long studied how wages and employment vary over the business cycle, and respond to international trade.

In recent decades, the subfield of personnel (organizational) economics arose. This area focuses on the firm: how does it design policies to effectively attract, develop, retain, and motivate employees? Its interests are closely aligned with those of traditional human resource and industrial relations researchers, but gives greater emphasis to theoretical concepts from microeconomics, as well as econometric empirical methods. This field has developed a large, robust body of research that complements work from related disciplines. Notably, Bengt Holmstrom shared the 2016 Nobel Prize in Economics for his fundamental work on the theory of performance evaluation and incentives.

At the 2016 Academy of Management Conference, a group of distinguished compensation researchers held a panel discussion on the future of compensation research. Their remarks were compiled into an article published in this issue. Soon after the panel, Charles Fay commissioned a similar discussion of by leading economists doing compensation research: Michael Gibbs (University of Chicago), Kevin Hallock (Cornell University), Edward Lazear (Stanford University), Kevin Murphy (USC), and Canice Prendergast (University of Chicago). The format is slightly different, as this group did not meet together. Instead, Michael Gibbs held a conversation with each separately and compiled their remarks. The panel consists of economists who have contributed to compensation research by developing theory, conducting empirical studies, and collecting new types of datasets to study new questions.

Question 1. What are the chief successes & failures in compensation research?

Canice Prendergast. Probably the great success is that we have a good understanding of trade-offs associated with different kinds of compensation packages. Those trade-offs can be in terms of level of compensation, so for example it gives us a lens through which to understand why people working in venture capital, or CEOs, etc. get paid more. We have a metric for understanding that, and it allows us to calibrate whether those numbers look reasonable, or whether they look terrible. An example would be the work that Marianne Bertrand has done with Sendhil Mullainathan using the techniques that we have to suggest that perhaps there is some rent seeking by CEOs. You need to have some framework to do that. I would say that's a great success.

The second success on that dimension is the literature on multitasking. One of the big issues we have to worry about when designing compensation is the multi-dimensionality of what people do. School teachers are the canonical example of this. We have lots of evidence on No Child Left Behind. We now see the extent to which paying school teachers on test scores changes their teaching. Not so much on "teaching to the test," because we don't have good data on that, but that it tends to focus the interest of school teachers on the marginal kid; the student that gets the teacher over the line for a bonus, rather than the average kid. That is an example of tradeoffs that we have learned a good deal about.

Another success, and I give credit to our new Nobel Laureate [Bengt Holmstrom] is the idea that there are interesting interactions between compensation and ownership. We had some idea about this from earlier work by Grossman and Hart, but I think one of the contributions that Bengt has made is to help us think about the relationship between compensation and other issues.

One of the other things the literature has been quite useful on is identifying some of the issues that arise in the context of team production. We can talk about bad things like free rider problems, peer pressure and so on. There's not a lot of evidence on this, but we at least conceptually know many of these things.

There are also a small number of papers that show empirically the extent to which pay for performance seems to work pretty well.

Kevin Murphy. One of the interesting things about executive pay research is that it has attracted researchers from myriad fields. You've got accounting, finance, labor economics, organizational economics, personnel economics, law, strategy and management, so there's richness in perspective that's unusual in the social sciences.

But of course, a failure is that researchers in these fields don't talk to each other very often, and when they do talk to each other, it often seems like they're talking in different languages or at cross purposes. There is a real missed opportunity here in cross disciplinary learning that is not being done.

Another failure in the literature is that executive compensation is complicated, and many researchers fail to grasp the consequences and importance of that complexity. For example, in a lot of research on executive pay someone takes a total compensation number pulled from a machine-readable source, and thinks of that the same way they would think of a wage in a standard wage regression. And yet, when you look at what's underneath that total compensation number, it has base salary, annual bonuses, discretionary bonuses, long term incentives, restricted stock, stock options, performance shares, retirement benefits, and perquisites ranging from health club memberships to personal use of the corporate jet. Many of these forms of compensation depend on how performance is measured over a single or multiple years. It's not obvious how or when to measure them.

Stock options, for example, typically have terms of up to 10 years. The question for the researcher is whether stock options should be counted as compensation when granted, or only when exercised? In addition, executives routinely receive lump sum amounts at various points in time. They get signing bonuses when they join their firm, severance payments when they leave, and change-in-control payments when their companies are taken over. For someone who is treating this just like a standard wage regression, it's not obvious how or when to measure those aspects.

Finally, different components of compensation impose different amounts of risk on executives. Payoffs from stock options are inherently riskier than the payoffs from restricted stock, which are in turn more risky than base salaries. We know that risk averse and undiversified executives will place a lower value on riskier forms of compensation, yet most studies of executive pay simply and blindly add all these different forms of compensation together, where many times they don't even know that they're adding. Fortunately, there are more and more sophisticated researchers, so the situation is improving.

Michael Gibbs. The theory of incentive compensation is a great success story. (An even larger literature in labor economics on other aspects of compensation exists, but I'll focus on incentives). We have a robust and rigorous framework for thinking about many practical issues and tradeoffs in measuring employee performance (especially numerically) and tying it to explicit and implicit rewards. This structure helps us understand how incentive compensation must account for employee risk, how it may distort behavior in multitask work environments, and how employees may game rewards by manipulating performance measures. It can accommodate individual and group work, and relative performance evaluation. This structure has had enormous empirical application in personnel economics, accounting, finance, governance, marketing and strategy. By and large, evidence accords well with the theory, if researchers apply adequately sophisticated use of available theoretical ideas. It also has tremendous practical value. For example, a growing literature uses these ideas to study the role of imperfect incentive contracts at banks in the financial crisis of 2007-8.

What may be a "failure" is that academics sometimes think from too-simple models while reality is more complicated. For example, many economists seem to believe that large firms use promotion tournaments. Doing so in practice would generate some serious tradeoffs (e.g., tying rewards to hierarchical structure rather than using other methods of incentive pay; motivating those who lose the contests; giving up the option of hiring from outside the firm). Tournament theory is quite valuable but does not translate directly into real world applications very often. Our theoretical ideas are complex and must be applied thoughtfully and appropriately to different contexts.

While our theories are very relevant, they do fail to capture practical limits on design and implementation of policies. Most incentive plans are very simple (or “bureaucratic”) in practice, using 1 or 2 performance measures and linear bonus formulas. Some theoretical justifications for linearity have been proposed. However, it does seem clear that the compensation, evaluation and incentive policies that firms use are far simpler than what theory suggests. At the extreme, it appears that many large firms impose identical evaluation and compensation policies on their entire workforces, and that such plans are quite simple in structure. “One size fits all” policies are difficult to explain with existing theory. Even more puzzling, perhaps, is that many firms seem to use firm-wide incentive schemes in which many or most employees share a large bonus pool, based on some aggregate performance measure such as annual profits. In a larger firm, such a performance measure is nearly impossible for any one employee to affect no matter how hard they try. Moreover, most of these schemes seem to give very small bonuses to a typical employee. It is difficult to see how such schemes would provide incentives. Documenting the extent of such practices is worthwhile. Providing better explanations for them would be even more useful.

Edward Lazear. Most of the theory is already well mapped out. You’re never done, and there are always new things one can think about, but it’s good at this point. We have a good sense of what the field is about on the theory level. I would say the biggest failure historically was our inability to get really good data. More recently that’s turning from a failure into a success. There have been a significant number of firm-based datasets that are giving us good insight. This research on incentives has been very productive.

Question 2. What are the most important open questions? Understudied topics?

Michael Gibbs. The largest open question, in my mind, is subjective performance evaluation. Essentially every job’s pay scheme makes use of subjectivity in some way. Most employees get subjective evaluations, formally or informally by supervisors. Many don’t get any quantified performance metrics, so their only evaluation is subjective. Even employees who do not receive subjective evaluations have some elements of compensation that are determined with judgment: the magnitude of raises and bonuses, promotions or job assignments, etc. Despite the evident empirical relevance of this topic, it has been almost completely ignored by economists. There is some literature on implicit and relational contracts, but it tends to be abstract or focused on questions of organizational boundaries. It would be nice to see theorists develop practical, personnel economics-style theory about subjective performance evaluation. It would be even more interesting to see empirical economists collect more and better data about various uses of discretion in performance evaluation and incentives.

For example, subjective evaluations can be used to reduce performance measure risk, if the supervisor is aware of exogenous factors affecting the employee's output, and can use judgment to partially filter them out of the final evaluation. Moreover, the supervisor can do this especially for negative effects on employee performance, which is probably optimal given that employees are risk averse. One of the most useful roles for subjectivity may be to balance multitask incentives, since performance may be easy to quantify for some tasks, but difficult to quantify for other tasks. A salesperson might get a commission on sales, but also a bonus based on supervisor judgement about the quality of customer service. Such evaluations are ex post, so they may also be used to punish the employee if manipulation of numeric performance measures is detected. Finally, they may be used to motivate the employee to use delegation over "controllable risk" (see below) more effectively.

Economists have largely ignored intrinsic motivation. I am not quite sure why, given its obvious importance (including in our own jobs). This is relevant since interactions between pay for performance and intrinsic motivation could be negative (as most psychologists argue), positive, or vary depending on how incentives are design. The latter is my view – incentive pay undermines intrinsic motivation only if the performance measures used distort the employee's efforts away from tasks we normally consider as being intrinsically motivated. It seems fair to assume that all employees are motivated in part by intrinsic considerations (the nature of the work they perform). If so, incentive compensation might be viewed as a complement to intrinsic motivation, and designed accordingly. The optimal performance evaluation (behaviors or outputs to reward or punish), and the intensity of incentives, should then vary with the nature of the job. This seems an important element that we have completely ignored. Once again, it would also be nice to see some attempt to empirically study these issues.

Finally, I'll briefly mention three other interesting and neglected topics. We have a poor understanding of group incentives. How important are free-rider effects? How large does the group need to be before such effects become significant? Are there practical ways to overcome them (since Holmstrom's "breaking the budget constraint" does not seem to be used much in practice)? Second, how do evaluation and incentives affect employee risk-taking or risk-management? Third, what contributions, if any, might behavioral economics rigorously add beyond the comprehensive theory and empirical insights already developed by personnel economists?

Kevin Hallock. From an empirical point of view, there is a great deal of work that needs to be done around the concept of total rewards, as opposed to wage and salary income. Nearly all of the research to date is based on information on wage and salary income. This is not because that is really the most important issue, but rather because the data are simply available. There are so many additional attributes of compensation packages (both monetary and non-monetary) that are under-studied.

Canice Prendergast. We have not learned much about are how well things other than pay for performance motivate people. By that I mean the form of compensation. Firms spend a lot of time thinking about giving people vacations, having prizes, and a whole set of other types of compensation that we don't understand well. We don't understand a lot about the *form* of compensation. Usually when we write down models, we translate all of these into some monetary equivalent. We're very good at analyzing the trade-off between pay for performance and salary, but that's all the same dimension.

We're not very good at understanding other ways in which we reward people. For example, do people care about titles? Yes. Do we know how much they care about titles? No. Firms seem to have some handle on imagining what that might be. Do we know how much people care about themselves relative to their peers? We know a little bit about this but not much. What brought this home for me was our new building. When we moved into this building we had to assign offices. Econ. 101 would tell us that you care only about the characteristics of your office. But overwhelmingly people cared about their office compared to everyone else's office. Understanding how firms are aware of this, and how they leverage this, would be useful to know.

The biggest two things that are lacking in the literature are a central set of questions that we want to try to answer; and more fundamentally, we are very much limited by the absence of good empirical facts. My view on this is that fields can only thrive if they have both a lot of facts, and a set of central questions to be answered. I think ultimately the literature on compensation has neither.

Think about the last 30 years of people doing work on compensation research. Any time a typical paper gets written down it has some formal model where they say that expected utility responds to some form of observed efforts that you carry out, and effort responds to compensation. After 30 years we hardly know whether that fact is true. We have maybe two or three papers that estimate this, in a literature that probably has two or three thousand theoretical papers about this idea. That's a fundamental limitation of the literature. We are still in a world where we can imagine whatever it is we feel like imagining. You can't do that in public finance. There are elasticities that we've estimated. You can't do that in economics of education because we have elasticities that we have estimated. You can do whatever theory you want in our field because we don't have good empirical facts.

One example is that the theoretical literature on compensation has moved on to imagining extremely complicated compensation contracts which we'll never see in the world. Another failing of the literature is that we are not trying to understand why the world is simpler than our models typically claim it's going to be. When fields don't become calibrated by facts they become untethered, and I fear that is what has happened to the literature on compensation. The real danger then is that it becomes an "academic" discipline – only of interest to the academics, not the practitioners.

Edward Lazear. I would say primarily selection and hiring. The way I think about it is that one of the most important aspects of making a firm productive is getting the right people. While incentives are important, I think of selection as being probably even more important, because if you recruit the right person, the incentive stuff largely takes care of itself. We don't have very good models or ways of thinking about recruiting; it's still largely a black box.

Kevin Murphy. Executive pay at this point is a pretty mature literature that's grown even faster than CEO pay itself. The low-hanging fruit appears periodically when there are new disclosure rules, but those opportunities disappear quickly and there are always people itching to be the first to take advantage of any new rule. Understudied topics are in two buckets: things that are just hard to study, or things that require data that aren't easily attainable using the standard machine-readable sources.

While we know a lot about the composition or determinants of pay, there's relatively little consensus on the consequences of pay, such as demonstrating causal links between the level and structure of pay, and subsequent performance. These links have been hard to establish because stock prices are relatively efficient, so we shouldn't expect researchers to be able to predict future stock price performance based on current publicly available data on pay. The links have also been difficult to establish because everything is endogenous. In corporate finance, the "endogeneity police" have taken over; labor economics is the same way.

Question 3. If you could refocus theorists or empiricists, how would you redirect those literatures?

Edward Lazear. I would say it's about right. Most of the work is empirical, and that's as it should be, because the earlier work was mostly theoretical. Initially we were unable to get good datasets, but we made good progress on the theory side. What's lacking is even more data and information on how things actually operate inside firms, so I would say that the current emphasis on empirical research is close to right. There are still some guys working on theory; e.g., me and Mike Waldman, but we try to couple it with empirical work.

With respect to what's missing in agency and related theory … again I go back to hiring models. What does it mean to get the right person? Who is the right person? How should we think about that problem? It's a matching issue, but thinking about matching at a much less abstract level than how it has been modeled to date.

I have a paper on this with Kathryn Shaw and Chris Stanton, forthcoming in the *Journal of Labor Economics*, but it too is somewhat abstract. What we're trying to do – and this is a more general theme – is to connect abstract agency literature with institutional features of the real world, but in a more formal way. Our notion is that job slots are important in practice. Except in search theory (which is very abstract) there's no notion that a slot matters. Let's say two similar rookie economists apply for a job at a top university, but there is only one available position. One will get a great job, while the other won't despite their very similar abilities. We see this happening in the academic market.

If you think about standard production theory, there's no notion of a slot per se. In my example, the two applicants should end up doing the same job, moving along the production frontier in a smooth, continuous way. A slot, though, is a discontinuous, discreet notion. Allocation of people to slots has significant implications. In our paper we find that where you end up depends in large part on luck; not just your luck, but also who else happened to be applying for a job at the same time. Cohort effects are for a whole group of people, but they have that same flavor. If you happen to be born into a cohort during a period of recession, it's bad luck beyond anything having to do with your skill set. The question is, why do such effects last over time?

That's an example of making models less abstract and more relevant to what actually goes on in the real world, but also thinking about the real-world issues in a formal way instead of a loose way. I'm hoping that more economists will do theory of this type.

Michael Gibbs. Above I mentioned subjective evaluation and intrinsic motivation as important topics that have been largely ignored. In addition to those, I would like to see both theorists broaden their concept of the employee's job. Almost all theory supposes that an employee exerts multiple types of "effort" that have marginal disutility for the employee, and generate "outputs" for the employer, but this is quite abstract. A very important dimension of the job is "knowledge": information, observations, experience, ideas, etc. that an employee possesses and might use on behalf of the firm (or on behalf of himself in manipulating performance measures). Firms give employees more or less discretion to use such knowledge in their work, depending on how valuable that knowledge may be to the firm, and the extent to which incentives can be aligned with firm objectives. Thus, employee learning and information developed on the job, extent of decentralization, method of evaluation, and strength of incentives are all endogenous.

Moreover, this view emphasizes that “risk” is not as simple as originally modelled by Holmstrom and others. A factory manager cannot control the weather, but he can manage how a snowstorm affects factory operations and profitability. Thus, risks are sometimes uncontrollable and should be filtered out of the employee’s evaluation as much as possible. However, other risks are controllable and should not be filtered out. The employee should be measured on such risks, and given incentives to manage them effectively. These basic ideas are somewhat understood (Canice has done important work in this area, for example). However, this key point is rarely recognized or given much attention in theoretical or empirical work.

Learning by the employee is also important for understanding intrinsic motivation. A key element of models of intrinsic motivation in organizational behavior is learning. When an employee is performing unfamiliar tasks and learning new skills, they are forced to think about what they are doing, which is perhaps the most important type of intrinsic job motivation. This suggests that a firm might design a job so that employees learn things that may be of value to the business (identify quality issues, solve problems, improve customer service, make continuous improvements). Intrinsic motivation may arise in situations where firms use more decentralization, leading to additional interactions between performance evaluation and incentives and other variables of interest.

New technology has had dramatic effects on job design, demand for various types of skills, and compensation for over 200 years. A large literature explores how the information technology revolution in the last three decades generated routine-biased technical change, and increased wage inequality. More recently, this literature has emphasized how technology is polarizing labor markets by automating middle skill jobs more than low and high skill jobs. These effects may be changing, as recent advances in artificial intelligence enable machines to perform more tasks normally considered low or high skilled (e.g., autonomous vehicles, some forms of medical diagnosis). Given the enormous importance of these developments, it would be of great interest to delve more deeply into understanding how compensation is affected by the composition of the job (which types of tasks are performed, and skills are required).

Kevin Hallock. I am an empiricist so let me concentrate my attention on that. I think empirical work needs to focus more on total pay. There is a massive empirical literature on wage gaps (e.g., the male-female wage gap) but nearly all of this is based on either hourly wages or annual salaries. To the extent that some groups have differentially better benefits, the results of those literatures may be misleading. For another example, there is a documented wage gap between individuals who have disabilities and those who do not. However, in recent work with Xin Jin and Linda Barrington, we hypothesized that the total pay gap between the two groups would be smaller in percentage terms than the wage gap, because those individuals with a disability may differentially select into positions with relatively better benefits. Using very unique data this is precisely what we found.

Kevin Murphy. The predominant theme in the CEO pay literature over the past 15 years has been asking whether CEO pay is set in an efficient managerial labor market, or as the outcome of managerial power (as emphasized by Lucian Bebchuk at Harvard). I think that debate is kind of tired. You can see why it's important, because issues of income inequality have been tied to executive pay. We're pretty tolerant of anybody who makes a lot of money because of hard work, innate ability, or better ideas. We're very intolerant of anyone who gets high pay because of perceived corruption. It thus becomes important whether or not executive pay reflects a competitive market for managerial talent, or managers passing high levels of pay through passive or complacent boards of directors. However, focusing on that distinction not only ignores the fact that they're not mutually exclusive, but also diverts attention from what might be even more interesting issues.

One important issue is that you can't explain CEO pay without understanding the role of the political sector. CEOs have the most highly regulated pay in the labor market. Many unique tax laws are written just for top executives of publicly traded firms. For example, since the mid-1990s the IRS has capped deductions companies can take for non-performance-related pay exceeding $1 million. This highly discriminatory rule only applies to publicly traded companies and only to the four or five executives whose pay is disclosed in annual proxy statements. In addition, since the 1980s, change-in-control payments perceived as "excessive" (which Congress arbitrarily defined as three times average W2 compensation over the prior five years) has been hit with both deductibility limits and excise taxes imposed on the executive; similarly excise taxes on executives were imposed in 2004 on in-the-money option grants and certain forms of deferred compensation.

In addition, since the 1930s companies in the U.S. have been required to disclose details of pay for individual executives. Why do we have disclosure rules for executives in publicly traded firms, when we don't see them anywhere else in the U.S., except for the public sector? The reason we have them for CEOs has nothing to do with shareholders wanting this information to make more informed decisions. It comes from a political sector that's uniquely focused on the level of pay, and they want to use disclosure as a tool to mitigate what they perceive as excesses in the level of pay. One of the interesting developments over the last 80 years is how little evidence there is that disclosure actually results in lower pay. There is a lot of evidence that disclosure can increase the level of pay. CEO pay became a political issue in the early years of the Great Depression, but has taken center stage since the Presidential election in 1992, when all major candidates (including Bill Clinton) took a critical stand on executive pay. When Hillary Clinton announced her candidacy for the 2016 election, she brought up what she saw as CEO excesses in her first five minutes.

Therefore, if I could redirect empirical work in CEO pay, it would be to understand and embrace the richness of pay, and to move away from the narrow debate on whether CEO pay is the result of efficient contracting or managerial power. On the theory side, too much of the literature is still directed by

Holmstrom’s Informativeness Principle from 1979 – that the way to think of CEO pay is as a hidden-action problem. The CEO takes an action and the board can’t see it. If they could see it, they would just tell the CEO what to do and get the first-best level of effort. Under the Informativeness Principle, the purpose of incentive-based pay isn’t to provide incentives, but rather “as if” we’re just taking signals of performance and figuring out whether you actually did what you were supposed to do.

Well, of course that’s a very simplistic way of think about what CEOs actually do. It’s also very simplistic to think that shareholders would know what they wanted CEOs to do. Even if they followed them around all day with a clipboard and knew exactly what they did, shareholders wouldn’t have any idea whether the executives took the right actions or not. We model this simplistically with hidden action problems without recognizing that the real decisions CEOs make are which project to take, what kind of business lines to be in, what kind of business lines to get out of, what the overall strategy of the company should be, and things that just aren’t translatable into a unidimensional single action. We do that for simplicity but sometimes we start taking it a little too seriously, too.

Canice Prendergast. I was a journal editor and had the right to decide what the world got to see: I would publish a lot more fact papers. Even fact papers that aren’t identified to the extent that our field requires. 20 years from now the world doesn’t need another extremely clever theory paper. What the world needs is a million facts that eventually answer a question we want to know the answer to. As a field maybe everybody should come together, have one of these “come to Jesus moments” that we can’t make progress unless we know more empirical facts. We don’t know much. That would be the starting point.

It would be very valuable if, over the next five years, a group of academics would do the best they could to answer the following question: to what extent do people change their behavior when you give them some formal reason to be more motivated? The most common formal reason to be more motivated is by using some form of pay for performance; equity, commissions, whatever. That would be one question I would love to know the answer to. Amazingly, after 30 years of work I do not know that answer. I know it for windshield installers [Lazear's Safelite study], but beyond that do I know it for anyone else? Not really. You can't have a field where you only know the answer for windshield installers.

Question 4. What new types or sources of data would you like to see collected?

Kevin Murphy. First of all, I'm not one who wants to encourage the SEC to impose any more disclosure rules. If you look at a typical proxy statement these days, it can easily run 60 to 100 pages, with the overwhelming majority dedicated to executive compensation. 50 or 60 years ago, the proxy statement would be a few pages with a brief description of compensation and maybe a table showing what it is. These days, whenever there is an apparent excess or abuse of executive compensation, a new disclosure rule is passed focused on that issue. We virtually never see rules that require less disclosure.

There would be some fruitful things to do if we had proprietary data on executives below the top five. We've spent so much time focusing on the CEO and the other people mentioned in the proxy statement. We sometimes treat the CEO as a sufficient statistic for compensation within the organization but of course it isn't. We don't know much about how the structure of CEO compensation affects subsequent performance. We know even less about the structure of pay below the CEO level, and the performance-consequences of that structure. What are the performance implications of broad-based incentive plans? How important are broad-based incentive plans?

I would love to see – but don't want its disclosure required – data on CEO wealth portfolios. A lot of our theories, for example about risk premiums executives should demand for accepting risky compensation, depend on the CEO's outside wealth or the percentage of wealth tied up in the firm. Interesting research could be done if we knew more about the portfolio of the CEOs beyond what's related to the company.

Canice Prendergast. I would like to know the empirical trade-offs associated with the most common way in which people's pay gets changed – which is that the boss decides what raise to give. Again, we have lots of theory papers and I've been as responsible as anyone for writing some of these. Empirically, what are the pros and cons associated with giving discretion to people? I think understanding this would be really valuable because it's a first-order issue. How does my pay get determined by the Dean? I get a letter every year that say's "This is your pay increase." Is that better than having something formal? I think we know it's usually the alternative to some mode of pay that's motivating something explicit, but I don't know how subjectivity in pay works, to be honest. I don't know how well it works. I don't know the circumstances in which it works.

One thing I think we know a bit about, probably more than some other things, is the impact of peers. To what extent do peers monitor us? To what extent do peers motivate us? Let's use academics as the example: what motivates us? Realistically, is my pay raise what motivates me? I don't think so. We all know within a reasonable band what we're going to get, maybe because of subjectivity issues. If we're great we're not going to get that much more if than if we're not. It's not pay that motivates us. Truthfully, it's how our peers perceive us, and that's because in many occupations we tend to be members of a profession and professional standing matters to us. How this effect interacts with compensation might be an interesting question.

It seems a bit weird when we write down models in which pay motivates us, but to a first approximation it doesn't motivate any of us. It's something else that gets us up in the morning; we like what we're doing.

The other thing that would be useful for me to know, outside of a very limited literature, is the interaction between compensation and intrinsic motivation. I mean in the field. I find much of the work on this question narrowly focused. A researcher might try to tease out the idea that if you pay someone to do something they enjoy it less. You can put people in a lab and get that kind of finding. Another researcher might ask, if you gave someone $8 million to make a 4-foot golf putt, might they get so nervous that they would perform worse than if you gave them $8 to make the putt. Those things are fine, but they're far removed from anything most of us do.

Let's take the following example. I like my job. Suppose there was a prize for a great paper. Would that make me more likely to write one? Less likely to write one? I know most of the literature is totally irrelevant for this issue. I know the idea of the $8 million putt and the $8 putt is totally irrelevant. The reason that affect occurs is physiological: I get an adrenaline rush when there is an $8 million putt. But in my job, I write one or two papers a year. I cannot have an adrenaline rush for 8 hours a day for 360 days a year. It's physically impossible. Those are useful types of papers about what happens when you get into a crisis, which is a rare, high stakes situation, but that's not the world we live in. If I want to know why like a guy who's never been a trader before might panic when he's a trader, fine, but it's totally irrelevant to what we usually do.

Edward Lazear. I would say that there are two paths. One is more matched worker-firm datasets of the kind that we now have from many countries in Europe. The Danish data, for example, are fabulous. It would be great to have data of this sort at all levels and in more countries. In the U.S., we are halfway there with Census LEHD [Longitudinal Employer-Household Dynamics] data. However, that database is difficult to gain access to, and is nowhere near as rich as data produced in other countries. For example, the Danish data provide nearly all relevant information about all citizens, matched not only with their firms, but also with firm performance.

The second path is much more information about individual worker output. That includes subjective evaluations if necessary, but also measures of worker output that are as close to objective as possible. An example is my Safelite study. It was trivial to measure output because you could count the number of windshields a worker installed. In most jobs, it's not that easy, but to the extent possible we should be thinking hard about how to measure individual output, and how to relate that to wages and other personnel outcomes.

Michael Gibbs. More data of any kind is always good. In the last 2-3 decades we have collected many new kinds of datasets with different variables, practices, and settings, but there is still a great deal that we do not know in detail. I would love to see empirical researchers try to measure qualitative issues with more rigor: employee outputs such as quality and service; inputs such as “effort,” diligence, and collaboration; intrinsic motivation; subjective evaluation; perceived trust of the manager; extent to which the manager coaches and trains, etc.

We know very little about organizational structures, or patterns and methods of decision making. We know little about innovation processes and outcomes inside firms. We rarely obtain information on which employees work together in groups (or work together occasionally, say in cross-functional teams), social networks within the firm, or who supervises whom. These are not compensation issues per se, but I’ll mention them anyway. There are a lot of important organizational issues that we rarely discuss, almost surely because we have little data on which to base such discussions.

The matched worker-firm datasets that have been generated recently are fantastic for providing rich detail usually only found in single-firm datasets, across many firms or even an entire economy. The next step is to collect firm financial data as well; some progress is already being made. Along similar lines, having detailed internal accounting data (e.g., at individual cost-center or organizational unit level), matched with detailed personnel records, might lead to some interesting research. Such accounting data might proxy for employee or group performance.

We have very little data on firm policies. If someone could collect information on various personnel policies across a large sample of firms, with reasonable details, it would be quite interesting. We do not have a good sense for what policies firms actually use, with what prevalence, in what combinations, and in what circumstances. A challenge, of course, would be to figure out how to measure practices consistently across a large set of firms. A second challenge would be that many policies are somewhat informal or implicit.

Kevin Hallock. I have a number of thoughts on this. First, more data that follows employees as they move from one organization to another would be really great. There are examples of this for some kinds of employees in the United States (e.g., executives through proxy disclosure, leaders of nonprofits through 990s, athletes and other public figures). There are some inroads on this front, but more information following the same people would be helpful for many reasons, not the least of which is that researchers can then control for the individual in important ways.

Second, in the absence of data that follows individuals from organization to organization, it would be great to know more about individual employees. We often know things like age, gender, race and education, but it would be great to know about disability status, the kinds of schooling individuals have had, the sorts of mentoring they have experienced, and personality and ability traits.

Third, I would love to have more information from more sources on total compensation. In the United States, we have pretty good data on wage and salary income, but not such great data on other important attributes of pay that make up total rewards. It would be great to know about the value of benefits, work environment, tools available on the job, and even about colleagues.

Question 5. What research methods would you like to see used more?

Kevin Hallock. I would love to see more research designs that would help practitioners and scholars feel more confident about whether there are *causal* effects of a particular change in a compensation practice (or HR practice more generally) on outcomes (say, performance, engagement, happiness, whatever). Too much work is based on correlations and is not convincing regarding the causal mechanisms. There is some feeling that managers or executives are indeed willing to change practices, but often don't have the time to see whether the practice led to a change in outcomes, or whether the change in outcomes might have been due something else that changed along the way. As a result, many organizations don't really know the true causal effects of their compensation or other HR practices. As a result, I wonder how much is lost in terms of lost engagement, productivity, profits etc. It could easily be the case that firms may be a lot more profitable and employees happier with a little bit more work in creating credible and objective ways to investigate the effects of HR practices on outcomes for employees and organizations.

Edward Lazear. I've been working in this field for many years, and I wish we were more like in the old days. What I mean by the old days is what I think of as Sherwin Rosen-style empirical economics, in theory but actually more on the empirical side. Sherwin's style was basically to use common sense and not cute strategies or overly complex models.

If you think about where modern econometrics has gone there are basically two areas. One is treatment effects methods, where you're looking at instrumental variables and trying to find cute identification strategies. The other version is structural modeling. My feeling is that both of those, while very clever, don't always give you the best information. I always thought of Sherwin as being a master at this, looking at data and making sense out of simple stuff, and knowing what he was really looking at. He would never fall into the trap of writing down a structural model, estimating the rate of return equaling 95%, then saying, "That's the number I get so that must be right." He would say that if the rate of return is 95% everybody in the world should be investing, so that cannot be right. Zvi Griliches was another guy like that – an unbelievable master at looking at data and really understanding what was going on.

That's a style that, unfortunately, has largely vanished. In part for good reason, because some weren't doing it in a thoughtful way. But I think that it's a lost art, and being a little bit more thoughtful and sensible about you're actually looking at is something we should try to teach our students.

Kevin Murphy. It's hard for me to talk about research methods when I'm pretty much a cross-sectional guy, or maybe even worse just a sample means guy, but clearly there is scope for finding more novel or compelling ways to address some of the endogeneity issues with pay. I'm not a member of the endogeneity police, but at the same time executive pay is a particularly challenging area. Everything is endogenous and there are very few natural experiments or obvious instruments. I am writing a paper on the role of compensation consultants, facing the challenge of finding an instrument that's related to the use of consultants but unrelated to the level of pay. It's pretty hard to do.

If a company has multiple divisions, you might be able to convince managers into conducting quasi-natural experiments and even randomization studies, in which divisions get different pay structures. But I haven't found a company that's willing to do that with its executive compensation policies.

It's also hard to find natural experiments. You can find some quasi-natural experiments that affect everybody at once, but there's not that many that affect selected groups of companies or individuals in order to run some of these tests.

Canice Prendergast. It would be useful for the field to decide what it would like to do; what it would like to know. For me, that would be a million facts. In research methodology, we need to change the bar that we typically use for publishing papers in journals, beyond that usual criterion that it has to be perfectly identified. One of the difficulties that social sciences have at the moment is that many care more about how precisely they estimate something than about what it is they're actually estimating.

That works in some cases. It works in public finance for the following reason: you can identify enough causal variables, like discount annuities or tax rates, or there's geographic dispersion. But there are fields where clear identification just doesn't work. It turns out that the two fields I happen to be in are good examples of this – traditional agency problems, and my new area of market design. If you change the market, you change the whole thing, so you don't usually have the sort of identification that you want. So usually it's just time series, but you know it's as good as it gets. Fields have to decide what's appropriate to them.

My sense is that we've made tremendous progress conceptually, but that at some point you go so far down the rabbit hole of conceptual advances that as a field you've got to pivot and say, "Just give me numbers." Even bad numbers are better than no numbers – eventually what will happen is that we'll converge.

Michael Gibbs. Personnel economics has had interesting "data entrepreneurship" for many years. I hope this trend continues with even more use of survey data (but with surveys written by economists so that the questions are better mapped to our concepts), and matched worker-firm-financial data. The current trend towards field experiments is quite useful in trying to get some handle on causality, since there are so many relevant variables affecting outcomes in organizations at the same time.

Firms are now doing “workforce analytics” – what we always called empirical personnel economics! They are digitizing quite interesting types of data, such as text used in performance evaluations, employee innovation, or social networks. We should try similar methods. They are also using “Big Data,” machine learning and neural network techniques to study employee and supervisor behaviors. I wonder if we may be able to use such methods to study our large personnel databases soon.

Finally, an interesting variation on a field experiment – more of a case study – would be to follow a large-scale organizational change over time, to see how it is implemented, what happens, how the organization evolves, etc. More generally, I would like to see more case studies of particularly interesting organizational settings. While these cannot be seen as representative or able to get a handle on causality, they have their own uses because the researcher can provide far richer data. Medical school students study pathologies, and perhaps so should organizational researchers.