

Kaplan, Steven N.

**Article**

## Executive compensation and corporate governance in the US: Perceptions, facts, and challenges

NBER Reporter

**Provided in Cooperation with:**

National Bureau of Economic Research (NBER), Cambridge, Mass.

*Suggested Citation:* Kaplan, Steven N. (2012) : Executive compensation and corporate governance in the US: Perceptions, facts, and challenges, NBER Reporter, National Bureau of Economic Research (NBER), Cambridge, MA, Iss. 3, pp. 1-6

This Version is available at:

<https://hdl.handle.net/10419/103228>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*

# NBER Reporter

NATIONAL BUREAU OF ECONOMIC RESEARCH

Reporter OnLine at: [www.nber.org/reporter](http://www.nber.org/reporter)

2012 Number 3



*Steven N. Kaplan*

---

## The 2012 Martin Feldstein Lecture

---

### Executive Compensation and Corporate Governance in the U.S.: Perceptions, Facts, and Challenges

Steven N. Kaplan\*

In this lecture, I explore some commonly held perceptions of executive compensation and corporate governance in the United States: 1) CEOs are overpaid and their pay keeps increasing; 2) CEOs are not paid for performance; and 3) corporate boards are not doing their jobs. For example, Bebchuk and Fried have concluded that, “flawed compensation arrangements have not been limited to a small number of ‘bad apples’; they have been widespread, persistent, and systemic.”<sup>1</sup> I consider the accuracy of these perceptions today, and discuss the implications and challenges that the evidence poses for researchers, boards, and shareholders.<sup>2</sup>

#### How is pay measured?

There are two ways to measure CEO pay. The first is estimated or grant-date pay. This includes the CEO’s salary, bonus, the value of restricted stock, and the estimated value of options issued that year. This is the compensation the board awards the CEO and, therefore, the appropriate measure for board governance effectiveness.

The second measure is realized pay. This includes the CEO’s salary, bonus, the value of restricted stock, and the value of options exercised that year. Because it uses actual option gains (not estimated val-

#### IN THIS ISSUE

The Martin Feldstein Lecture	1
Research Summaries	
The Production of Scientific Ideas	7
... Profitability of Currency-Trading Strategies	10
Work-Family Balance	14
Race, Income, and Political Efficacy	17
NBER Profiles	20
Conferences	23
NBER News	31
Program and Working Group Meetings	32
Bureau Books	34

---

\* This is a written and abbreviated version of the Martin Feldstein Lecture given on July 10, 2012. Kaplan is an NBER Research Associate and the Neubauer Family Distinguished Service Professor of Entrepreneurship and Finance, University of Chicago Booth School of Business. He also serves on public company and mutual fund boards. Douglas Baird, Effi Benmelech, Carola Frydman, Austan Goolsbee, Jeff Miron, Raghu Rajan, Amir Sufi, Luke Taylor and Rob Vishny provided helpful comments on this article.

# NBER Reporter

The National Bureau of Economic Research is a private, nonprofit research organization founded in 1920 and devoted to objective quantitative analysis of the American economy. Its officers and board of directors are:

President and Chief Executive Officer — *James M. Poterba*  
Controller — *Kelly Horak*

## BOARD OF DIRECTORS

Chairman — *Kathleen B. Cooper*  
Vice Chairman — *Martin B. Zimmerman*  
Treasurer — *Robert Mednick*

## DIRECTORS AT LARGE

Peter Aldrich	Mohamed El-Erian	Michael H. Moskow
Elizabeth E. Bailey	Linda Ewing	Alicia H. Munnell
John Herron Biggs	Jacob A. Frenkel	Robert T. Parry
John S. Clarkeson	Judith M. Gueron	James M. Poterba
Don R. Conlan	Robert S. Hamada	John S. Reed
Kathleen B. Cooper	Peter Blair Henry	Marina v. N. Whitman
Charles H. Dallara	Karen N. Horn	Martin B. Zimmerman
George C. Eads	John Lipsky	
Jessica P. Einhorn	Laurence H. Meyer	

## DIRECTORS BY UNIVERSITY APPOINTMENT

George Akerlof, <i>California, Berkeley</i>	Bruce Hansen, <i>Wisconsin</i>
Jagdish W. Bhagwati, <i>Columbia</i>	Marjorie B. McElroy, <i>Duke</i>
Timothy Bresnahan, <i>Stanford</i>	Joel Mokyr, <i>Northwestern</i>
Alan V. Deardorff, <i>Michigan</i>	Andrew Postlewaite, <i>Pennsylvania</i>
Ray C. Fair, <i>Yale</i>	Uwe E. Reinhardt, <i>Princeton</i>
Franklin Fisher, <i>MIT</i>	Craig Swan, <i>Minnesota</i>
John P. Gould, <i>Chicago</i>	David B. Yoffie, <i>Harvard</i>
Mark Grinblatt, <i>California, Los Angeles</i>	

## DIRECTORS BY APPOINTMENT OF OTHER ORGANIZATIONS

Jean Paul Chavas, *Agricultural and Applied Economics Association*  
Martin Gruber, *American Finance Association*  
Ellen Hughes-Cromwick, *National Association for Business Economics*  
Christopher Carroll, *American Statistical Association*  
Thea Lee, *American Federation of Labor and Congress of Industrial Organizations*  
William W. Lewis, *Committee for Economic Development*  
Robert Mednick, *American Institute of Certified Public Accountants*  
Alan L. Olmstead, *Economic History Association*  
John J. Siegfried, *American Economic Association*  
Gregor W. Smith, *Canadian Economics Association*  
Bart van Ark, *The Conference Board*

The NBER depends on funding from individuals, corporations, and private foundations to maintain its independence and its flexibility in choosing its research activities. Inquiries concerning contributions may be addressed to James M. Poterba, President & CEO, NBER 1050 Massachusetts Avenue, Cambridge, MA 02138-5398. All contributions to the NBER are tax deductible.

The *Reporter* is issued for informational purposes and has not been reviewed by the Board of Directors of the NBER. It is not copyrighted and can be freely reproduced with appropriate attribution of source. Please provide the NBER's Public Information Department with copies of anything reproduced.

Requests for subscriptions, changes of address, and cancellations should be sent to *Reporter*, National Bureau of Economic Research, Inc., 1050 Massachusetts Avenue, Cambridge, MA 02138-5398. Please include the current mailing label.

ues), this better measures what the CEO actually takes home. Accordingly, realized pay is appropriate for considering whether CEOs are paid for firm performance.

## Facts about pay

Using estimated pay, I look at data from 1993 to 2010 for S&P 500 companies (from S&P's ExecuComp database). What has happened to average estimated CEO pay (adjusted for inflation) since 2000? Most audiences believe it has increased substantially. In fact, Figure 1 (on page 3) shows that while average CEO pay increased markedly from 1993 to 2000, it declined by over 46 percent from 2000 to 2010. Median CEO pay also increased from 1993 to 2000, but has since declined. The convergence between the means and medians suggests that boards have become less likely to award large pay packages since 2000.

There are still some outliers that receive attention and likely drive the perception that pay has increased. For example, three CEOs received over \$50 million in estimated pay in 2010. The means and medians indicate that these are outliers and not the general rule.

ExecuComp also follows the CEOs of over 1,000 smaller companies not in the S&P 500. Average estimated pay for these CEOs, like S&P 500 CEOs, increased in the 1990s and declined in the 2000s. Today's average pay roughly equals its 1998 level.

Overall, then, estimated CEO pay — what boards expect to pay their CEOs — peaked around 2000, both for S&P 500 and non-S&P 500 CEOs. Since then, average estimated CEO pay has declined, returning roughly to its 1998 level.


While average pay has declined since 2000, it remains very high in absolute terms. In 2010, the average S&P 500 CEO received estimated pay of just over \$10 million. This is roughly 200 times the median household income in the United States and undoubtedly also contributes to the perception that CEOs are overpaid.

## Turnover

The average lengths of CEO tenures today are shorter than in the past. As a result, comparing CEO pay in the 2000s to CEO pay in the 1990s (and earlier) is not an apples-to-apples comparison. In the 1970s, 1980s, and

**Figure 1**

Average & Median Total Pay (estimated)  
of S&P 500 CEOs from 1993 to 2010 (in millions of 2010 \$)



mid-1990s, roughly 10 percent of large U.S. company CEOs turned over each year, not counting takeovers.<sup>3</sup> Since 1998, annual turnover has increased to an average of 12 percent, implying a decline in CEO tenure from ten to eight years. Including takeovers, tenures have declined from roughly eight years before 1998 to only six years since.

The decline in tenure implies that the CEO's job has become riskier over time. The shorter expected tenure arguably offsets roughly 20 percent of the increase in CEO pay since the early 1990s.<sup>4</sup> The true increase in CEO pay since then is lower than the compensation figures alone would suggest.

### How does CEO pay compare to that of other highly paid people?

Gabaix and Landier<sup>5</sup> argue that market forces can explain the increases in CEO pay. Using a simple competitive model, they show that CEO pay will rise as firms become larger because larger average firm size increases the returns to hiring more productive CEOs. They find empirically that the increase in CEO pay since 1980 can be fully attributed to the increase in large company market values.

Gabaix and Landier and others<sup>6</sup> focus on the market for public company top executives. But the same people also can become executives at private companies, become (or remain) consultants,

and, earlier in their careers, become lawyers, investment bankers, or investors. In a competitive market, similarly talented individuals should have done as well as CEOs over the last twenty or thirty years. The large increase in the share of pre-tax income earned by very high earners over that period, documented by Piketty and Saez,<sup>7</sup> suggests that this is plausible.

Accordingly, I compare the average estimated pay of S&P 500 CEOs to the average adjusted gross income (AGI) of taxpayers in the top 0.1 percent of the income distribution.<sup>8</sup> Figure 2 shows that average estimated pay for S&P 500 CEOs, relative to the average income of the top 0.1 percent, is about the same in 2010 as it was in 1994. S&P 500 CEOs

have seen little change in their estimated pay relative to other high earners since the early 1990s. And non-S&P 500 CEOs are worse off relative to the top 0.1 percent than they were in the early 1990s.


Over the last twenty years, then, public company CEO pay relative to the top 0.1 percent has remained relatively constant or declined. These patterns are consistent with a competitive market for talent. They are less consistent with managerial power. Other top income groups, not subject to managerial power forces, have seen similar growth in pay.

### What about the longer-term?

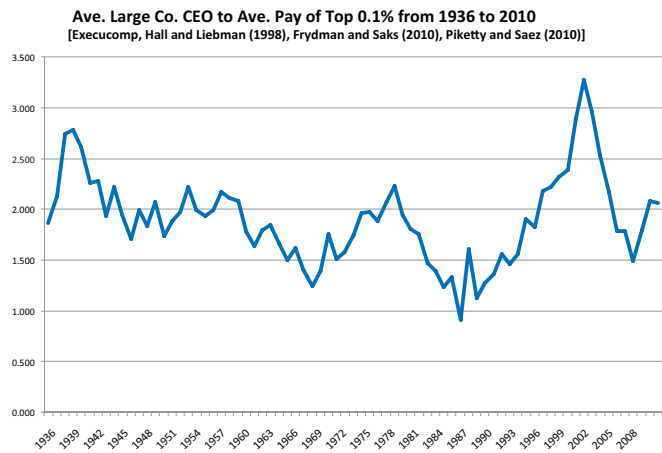
What has happened over the longer-term, since the 1930s? I staple together three data sets of estimated pay—ExecuComp data for S&P 500 CEOs from 1992 to 2010, the Hall and Leibman<sup>9</sup> data for large company CEOs from 1980 to 1992, and the Frydman and Saks data for large company CEOs from 1936 to 1980.<sup>10</sup> Figure 3 (on page 4) compares this series with the average AGI of the top 0.1 percent. Over the long-term, estimated CEO pay relative to pay of the top 0.1 percent has remained stable, averaging roughly 1.9. The ratio is particularly low in the 1980s, becomes unusually high in the late 1990s, and returns near to its long-term average in 2010. The unanswered question from these patterns is what drives the fluctuations.

**Figure 2**

Average Pay (Estimated) of S&P 500 CEOs to  
Average AGI of Top 0.1% of Taxpayers from 1993 to 2010



**Figure 3**



Sources: S&P 500 CEO Pay 1992 to 2010 from ExecuComp, Large Company CEO Pay 1980 to 1992 from Hall and Leibman (1998) Large Company CEO Pay from 1936 to 1980, Frydman and Saks (2010), Piketty and Saez (2010)

Figure 4 (below) shows the ratio of average estimated CEO pay to the average market value of the top 500 publicly traded companies (multiplied by 1,000). CEO pay was a higher fraction of market value in the 1930s through the 1950s than it was after 1960. Since 1960, however, the ratio has remained more stable, averaging 0.042 percent of market value. The ratio in 2010 was 0.036 percent. Since 1960, then, the data support the simple Gabaix and Landier story of a competitive market for talent. The unanswered question is why the pattern is so different before 1960.

Taken together, these long-run patterns suggest that a combination of the market for talent and firm scale have been meaningfully associated with large company CEO pay over a long period of time.

### Other specific groups

The previous analyses compare public company CEOs to those in the top income brackets. But public company CEO pay also can be compared to the pay of specific groups in those brackets that have similar opportunities or talents, particularly non-public company executives, lawyers, and investors.

Bakija, Cole, and Heim<sup>11</sup> study IRS tax return data between 1979 and 2005. They try to compare public and

private company executives by distinguishing those who receive the majority of their income in salary and wages from those who receive the majority from self-employment. The former are more likely to include public company executives; the latter, executives of closely-held businesses.

The pay of closely-held firm executives increased by more than the pay of salaried executives from 1979 to 2005. Closely-held firm executives also increased their representation in the top 0.1 percent, increasing from 9 percent

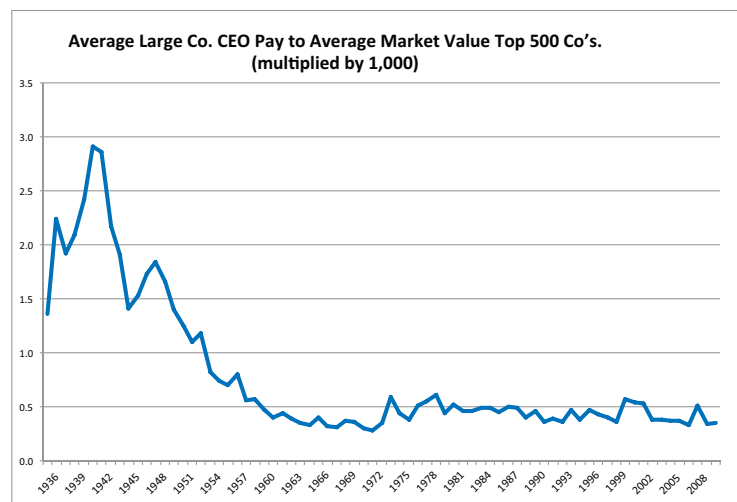
in 1979 to 22 percent of the top 0.1 percent in 2005. Over the same period, the representation of salaried executives declined from 38 to 20 percent.

Public company executives, those who should be more subject to managerial power problems, saw their pay increase less than executives of closely-held company businesses which are, by definition, controlled by large shareholders or the executives, and are subject to limited agency problems. This is notable because many of the salaried and closely-held executives likely come from the same executive pool and, presumably, can move between public and private company employment.

### What does this mean?

The point of these comparisons is to confirm that while public company CEOs earn a great deal, they are not unique. Other groups with similar backgrounds and talents—private company executives (as well as corporate lawyers, investors and others)—have seen significant pay increases where there is a competitive market for talent and no managerial power problems exist. If one uses evidence of higher CEO pay as evidence of managerial power, one must also explain why these other groups have had a similar or higher growth in pay. Instead, it seems

**Figure 4**



Sources: S&P 500 CEO Pay 1992 to 2010 from ExecuComp, Large Company CEO Pay from 1936 to 1980, Frydman and Saks (2010). Average market value of top 500 companies from CRSP from 1936 to 2010.



more likely that market forces have driven a meaningful portion of the increase in public company CEO pay.

Josh Rauh and I concluded that some combination of changes in technology, along with an increase in the scale of enterprises and finance, have allowed more talented or fortunate people to increase their productivity relative to others. This seems relevant for the pay increases of lawyers and investors (technology allows them to acquire information and trade large amounts more efficiently) as well as CEOs (technology allows them to manage very large global organizations).<sup>12</sup>

### Pay for Performance

Do CEOs who perform better earn more in realized pay — which includes option exercises and thus better measures what the CEO actually takes home? For each year from 1999 to 2004, Rauh and I took the firms in the ExecuComp database and sorted them into five size-groups. Within each size-group for each year, we sorted the CEOs into five groups based on realized pay. We then looked at how the stocks of each group performed relative to their industry over the previous three years.

We found that realized compensation was highly related to firm stock performance. In every size group, firms with CEOs in the top quintile of realized pay were in the top performing quintile; firms with CEOs in the bottom quintile of realized pay were in the worst performing quintile.

Frydman and Saks study the correlation between an executive's wealth and firm performance. They find that CEO wealth has been strongly tied to firm performance since the 1930s, and that relationship "strengthened considerably" after the mid-1980s.

The evidence, then, is consistent with realized CEO pay and CEO wealth being strongly tied to firm performance. The more difficult question is how much pay-for-performance is optimal, and whether current practices can become more efficient. Some argue that pay-for-performance is too low and should be increased.

Others argue that some pay-for-performance incentives, particularly in financial services, are too high.

### Are CEOs fired for poor performance?

CEO turnover levels have increased since the late 1990s, so CEOs can expect to be CEOs for less time than in the past. CEO turnover also has become increasingly related to poor firm stock performance.<sup>13</sup> This suggests that boards and the corporate governance system have performed better in their monitoring role since the 1990s.

Jenter and Llewellyn<sup>14</sup> present additional evidence consistent with this. They look at CEO turnover in ExecuComp firms from 1992 to 2004 and find "that boards aggressively fire CEOs for poor industry-adjusted performance, and that the turnover-performance sensitivity increases substantially with higher quality boards." In the first five years of their tenure, CEOs who perform in the bottom quintile relative to their industry are 42 percent more likely to depart than top quintile CEOs. This spread increases to more than 70 percent for firms with higher quality boards — more independent boards with greater stock ownership. As with pay-for-performance, the more difficult question is whether these differential departure rates are optimal and whether current practices can be improved.<sup>15</sup>

### What do shareholders think?

It would be useful to know what shareholders think. Fortunately, the Dodd-Frank Act of 2010 mandated that most publicly traded-firms hold Say-on-Pay votes — non-binding shareholder votes on the compensation of their top five executives. Say-on-Pay supporters believed that the votes would reduce the perceived CEO pay spiral and would increase pay for performance. Under the alternative view that pay levels and pay-for-performance are largely determined in a competitive market, the Say-on-Pay votes would be a non-event.

The law went into effect in 2011. The votes were overwhelmingly in favor of existing pay policies: roughly 98 percent of companies received majority support of their shareholders; more than 73 percent of companies received a favorable vote above 90 percent.<sup>16</sup> The 2012 votes have followed a qualitatively similar pattern. The positive shareholder votes for most companies seem inconsistent with top executive pay being driven largely by managerial power. Rather, the votes are consistent with a more market-based view.

### How have U.S. public companies performed?

Given the negative perceptions of CEO pay and corporate governance, one would think that corporate performance has been poor. The U.S. economy has gone through a financial crisis and recession, and the S&P 500 has declined from a peak of 1576 in 2007 to roughly 1400 today (August 2012). At the same time, CEO pay has declined. What has happened to operating performance?

S&P 500 companies have weathered the downturn surprisingly well. Median operating margins (EBITDA to Sales) increased from 1993 to 2007 and increased again, to their highest level in the period, from 2007 to 2011.<sup>17</sup> The National Income and Product Accounts, while they include public and private companies, also show that corporate profits as a fraction of GDP are at historically high levels. On average, then, particularly for non-financial companies, average operating performance has improved while average compensation has declined.

### Summary

To summarize, I have considered the evidence for three common perceptions of U.S. corporate governance. The evidence is somewhat different from those perceptions. For example, while average CEO pay increased substantially through the 1990s, it has since declined. Indeed, CEO pay levels relative to other highly paid groups today are comparable both to their average level in the early 1990s

and to their average level since the 1930s. And, the ratio of large company CEO pay to firm market value has remained roughly constant since 1960.

Furthermore, CEOs are typically paid for performance and penalized for poor performance. Finally, boards do monitor CEOs, and that monitoring appears to have increased over time. CEO tenures in the 2000s are lower than in the 1980s and 1990s, and CEO turnover is tied to poor stock performance.

In his 2012 work, Murphy concludes that executive compensation is affected by the interaction of a competitive market for talent, managerial power, and political factors. That conclusion is hard to disagree with. There have been corporate governance failures and pay outliers where managerial power surely has been exercised. And, CEO pay today is still extremely high relative to typical household income. At the same time, a meaningful part of CEO pay appears to have been driven by the market for talent. In recent decades, CEO pay is likely to have been affected by the same forces of technology and scale that have led to the general increase in incomes at the very top.

---

<sup>1</sup> L. Bebchuk and J. Fried, *Pay without Performance: The Unfulfilled Promise of Executive Compensation*, Harvard University Press, (2006).

<sup>2</sup> My talk addresses these perceptions. It is based on the more detailed treatment in my *Executive Compensation and Corporate Governance in the U.S.: Perceptions, Facts and Challenges*, *Cato Papers in Public Policy*, forthcoming. Two excellent and recent surveys, “Executive Compensation: Where we are, and how

we got there”, Kevin J. Murphy, forthcoming in *Handbook of the Economics of Finance*; and G. Constantinides, M. Harris, and R. Stulz, eds., and “CEO Compensation”, C. Frydman and D. Jenter, *Annual Review of Financial Economics*, 2 (2010), pp. 75–102, provide broader analyses and summaries of corporate governance issues.

<sup>3</sup> See K. J. Murphy and J. Zábajník, *Managerial Capital and the Market for CEOs*, 2008, and S. Kaplan and B. Minton, “How has CEO Turnover Changed?” *International Review of Finance*, 12 (2012), pp. 57–87.

<sup>4</sup> F. Peters and A. Wagner, “The Executive Turnover Risk Premium”, working paper, Swiss Finance Institute, 2012, which estimates this explicitly and finds that a 1 percent increase in turnover risk is associated with a 10 percent increase in pay.

<sup>5</sup> X. Gabaix and A. Landier, “Why Has CEO Pay Increased So Much?” NBER Working Paper No. 12365, July 2006, and *Quarterly Journal of Economics*, 123(1) (2008), pp. 49–100.

<sup>6</sup> See C. Frydman and R. E. Saks, “Executive Compensation: A New View from a Long-Term Perspective, 1936–2005”, *Review of Financial Studies*, 23, (2010), pp. 2099–138, and Murphy and Zábajník, 2008, *op. cit.*

<sup>7</sup> T. Piketty and E. Saez, “Income Inequality in the United States, 1913–1998”, NBER Working Paper No. 8467, September 2001, and *Quarterly Journal of Economics*, 118 (2003), pp. 1–39. (Tables and Figures updated to 2010 in Excel format, March 2012).

<sup>8</sup> This updates and expands the analysis in S. Kaplan and J. Raub, “Wall Street and Main Street: What Contributes to

the Rise in the Highest Incomes?” NBER Working Paper No. 13270, July 2007, and *Review of Financial Studies*, 23 (2010), pp. 1004–50.

<sup>9</sup> B. Hall and J. Liebman, “Are CEOs Really Paid Like Bureaucrats?” NBER Working Paper No. 6213, October 1997, and *Quarterly Journal of Economics*, 113 (1998), pp. 653–91.

<sup>10</sup> I thank Carola Frydman for providing them.

<sup>11</sup> J. Bakija, A. Cole, and B. Heim, “Jobs and Income Growth of Top Earners and the Causes of Changing Income Inequality: Evidence from U.S. Tax Return Data”, Working paper, Indiana University, 2012.

<sup>12</sup> See J.A. Parker and A. Vissing-Jørgensen, “The Increase in Income Cyclicity of High-Income Households and its Relation to the Rise in Top Income Shares”, *Brookings Papers on Economic Activity*, Fall, 2010, pp. 1–70, for a concurring view.

<sup>13</sup> See Kaplan and Minton, 2012, *op. cit.*

<sup>14</sup> D. Jenter and K. Lewellen, “Performance Induced CEO turnover”, Working Paper, Stanford University 2010.

<sup>15</sup> See L.A. Taylor, “Why are CEOs Rarely Fired? Evidence from Structural Estimation”, *Journal of Finance*, 65(6) (2010), pp. 2051–87, for an attempt at estimating this.

<sup>16</sup> See S. Mishra, “Parsing the Vote: CEO Pay Characteristics Relative to Shareholder Dissent”, working paper, Institutional Shareholder Services, 2012.

<sup>17</sup> See also “For Big Companies, Life Is Good,” Scott Thurm, *Wall Street Journal*, April 8, 2012.

# The Production of Scientific Ideas

Pierre Azoulay and Joshua Graff-Zivin\*

There is considerable evidence that the advancement of science influences productivity in the private sector of the economy. Thus, policymakers typically believe that public investments in science are important for long-run economic growth. But how do new scientific ideas come about? Apocryphal stories of Archimedes' eureka moment, or Newton's otherworldly contemplation interrupted by the fall of an apple, would have us believe that luck is an essential feature. Of course, if luck is all that is necessary to produce breakthroughs, then there is little room for scholarship on the subject. If, on the other hand, scientific knowledge production depends upon individuals, institutions, and incentives, then economic research should play an important role in increasing our understanding in this area.

While the pioneering work of Zvi Griliches, the founding director of the NBER's Productivity Program, set the stage for hundreds of subsequent empirical studies examining the diffusion of various technologies, comparatively little work has focused on the creation of the original technologies in the first place. This dearth of applied research on idea creation has not been for lack of interest but rather principally because of data limitations. As recently as 15 years ago, very little data were available to systematically study the scientific enterprise. Today, economists have at their disposal vast quantities of new data that allow them to link mentors and trainees, collaborators, and intellectual peers to char-

acterize the production team. The data on papers, patents, and citations enable one to trace out the impact of individual bits of knowledge as they are incorporated into the research activities of other research teams, as well as within private sector firms. Together with methodological advances in the analysis of quasi-experimental data, we have begun to credibly characterize this production process, the conditions under which scientists collaborate to create new knowledge, and the benefits that follow.

One important theme that has emerged from the recent literature is the notion that an increased burden of knowledge because of an ever-expanding scientific frontier has led to greater scientific specialization, longer training periods, and to an increased propensity to collaborate.<sup>1</sup> This realization has cast a pall over the potential for ideas-based growth, because it implies that innovation is becoming more difficult over time.<sup>2</sup> Ultimately, whether this pessimism is warranted is an empirical question, which has led us to explore in more detail the impacts of interactions among scientists for the pace of scientific advance, and whether these interactions occur because of geographic proximity, shared intellectual interests, or social connections.

## The Impact of Superstar Scientists

While the most important scientific work is much more likely to be produced as part of a collaboration than was the case only 40 years ago,<sup>3</sup> our own work suggests that the central members of these teams—whom we call “super-

stars”—continue to play an important role in shaping the rate and direction of scientific advance. Over the past ten years, we have gathered biographical information for a sample of 12,000 elite, academic life scientists, and combined these with precise measures of inputs (namely grants from the National Institutes of Health), outputs (publications and patents), and impact (citations to both publications and patents). Furthermore, we have linked these superstars to a much larger population of 200,000 academic life scientists in the United States, corresponding to most of the profession from the immediate post-war era to the end of the previous decade. Thanks to open-source software tools we designed for this purpose, we are able to locate all of these scientists in geographic space, identify their ties through co-authorship and citation networks, and assess the extent to which they work on similar topics.

Our first study in this area focuses on the benefits of exposure to superstar talent derived from formal collaboration.<sup>4</sup> The formation of collaborative teams is the result of a purposeful matching process, making it difficult to uncover the causal effect of collaboration on follow-on individual performance. To overcome the endogeneity of the collaboration decision, we use the quasi-experimental variation in the structure of co-authorship networks induced by the premature and sudden death of active superstar scientists. Our sample comprises 122 of these unfortunate events, and provides a unique opportunity to estimate the impact of the prominent members of scientific teams on their less-heralded collaborators when they work on other projects, as well as to

---

\* Azoulay and Graff Zivin are NBER Research Associates. Their Profiles appear later in this issue.



probe the mechanisms that undergird this influence. Our results show that upon losing a superstar collaborator in this way, scientists experience a long-lasting productivity decline, with a loss of 5 to 10 percent of their previous quality-adjusted publication output. Given the rich data we have gathered, we are also able to examine several competing explanations for this effect.

One view of the academic reward system provides the backdrop for a broad class of stories with a common thread: that collaborating with superstars deepens social connections, possibly making researchers more productive in ways that have little to do with scientific knowledge, such as connecting coauthors to funding resources, editorial goodwill, or potential coauthors. Yet, we find no differential impact on coauthors of stars well-connected to the NIH funding apparatus, on coauthors of stars more central in the collaboration network, or on former trainees. These findings do not jibe with explanations stressing the gate keeping role of eminent scientists.

Rather, the effects of superstar extinction appear to be driven by the loss of an irreplaceable source of ideas. We find that coauthors close to the star in intellectual space experience a sharper decline in output than coauthors who work on less related topics. Furthermore, the collaborators of stars whose work was heavily cited at the time of their death also undergo steeper decreases than collaborators of superstars who were less well known. Together, these results paint a picture of an invisible college of coauthors bound together by interests in a fairly specific scientific area, which suffers a permanent and reverberating intellectual loss when it loses its star.

This first paper focused on the effects of exposure to superstar talent through collaboration, but our second effort highlights geographic co-location as the channel of influence. We use as a source of variation the job transitions between distant institutions in our sample of elite scientists.<sup>5</sup> There again, the challenges involved in establishing causality loom large, since scientists might choose to switch jobs at

least in part based on the prospects of deeper interactions with colleagues or firms. We use a novel identification strategy that exploits labor mobility in a sample of 9,483 elite academic life scientists to examine impacts on the citation trajectories associated with individual articles (resp. patents) published (resp. granted) *before* the scientist moved to a new institution. This longitudinal contrast purges our estimates of most sources of omitted variable bias that can plague cross-sectional comparisons. However, the timing of mobility itself could be endogenous. To address this concern, we pair each moving scientist/article dyad (resp. scientist/patent dyad) with a carefully chosen control article or patent associated with a scientist who does not move to a new position. In addition to providing a very close match based on time-invariant characteristics, these controls share very similar citation trends prior to the mobility event. By analyzing the data at the matched-pair level of analysis, this simple difference-in-difference framework provides a flexible and non-parametric methodology for evaluating the effects of labor mobility on knowledge flows. Indeed, conditional on the assumption that the matching algorithm we employ successfully pairs articles and patents of comparable quality, we are able to present the findings in a straightforward, graphical form.

The results reveal a multifaceted story. We find that article-to-article citations from the scientists' origin location are barely affected by their departure. In contrast, article-to-patent citations, and especially patent-to-patent citations, decline at the origin location following a superstar's departure, suggesting that spillovers from academia to industry are not completely disembodied. We also find that article-to-article citations from scientists' destination locations markedly increase after they move. To the extent that academic scientists do not internalize the effect of their location decisions on the circulation of ideas, our results raise the intriguing possibility that barriers to labor mobility in academic science limit the recombination of individual bits of knowledge, resulting in a diminished rate of scientific exploration.

We are currently extending our research in this area along two dimensions. The "superstar extinction" study left open the question of whether we would observe the same negative impact on non-coauthors working in the same field. Our next study aims to ascertain whether co-authorship is required in order to be a full-fledged member of the invisible college of scientists formed around the star while alive. This question is important because it provides a lens through which we can examine whether scientific ideas are accessible to any trained scientist—exemplars of the pure public goods described in our economic textbooks—or whether they should be thought of as at least partially excludable—as would be the case if career success all but required direct connections with the scientific elite.

Second, we revisit our earlier work on the effect of elite scientist mobility, but with a focus on the creation of new knowledge as opposed to the diffusion of preexisting knowledge. Using a novel software tool, we can identify the peers of individual scientists based solely on their shared intellectual interests as indicated by keywords that tag publications—without any reference to linkages through co-authorship or citation. This opens the door to a deeper understanding of the process through which scientists position themselves in "intellectual space," and to the evolution of scientific fields over time.

## Incentives for Scientific Exploration

A distinct but related part of our innovation research agenda is how scientists choose projects, and the extent to which funding systems shape these choices. In collaboration with Gustavo Manso from the University of California at Berkeley,<sup>6</sup> we document that the features of incentive schemes embodied in the design of research contracts exert a profound influence on the subsequent development of breakthrough ideas.

In particular, we study the careers of researchers who can be funded through two very distinct mechanisms: investigator-initiated R01 grants from the NIH,

or support from the Howard Hughes Medical Institute (HHMI). HHMI, a non-profit medical research organization, plays a powerful role in advancing biomedical research and science education in the United States. It has also adopted practices that should provide strong incentives for breakthrough scientific discoveries: the award cycles are long (five years, and typically renewed at least once); the review process provides detailed, high-quality feedback to the researcher; and the program selects “people, not projects,” which allows for the quick reallocation of resources to new approaches when the initial ones are not fruitful. This stands in sharp contrast with the incentives offered to life scientists funded by the NIH. The typical R01 grant cycle lasts only three-to-five years, and renewal is not very forgiving of failure. Feedback on performance is limited in its depth and projects are funded based on clearly defined deliverables.

The contrast between the HHMI and NIH grant mechanisms naturally leads to the question of which incentives result in a higher rate of production of particularly valuable ideas. In the absence of a plausible source of exogenous variation for HHMI appointment, we estimate the treatment effect of the program by contrasting HHMI-funded scientists’ output with that of a carefully matched group of NIH-funded scientists who focus their research on the same subfields of the life sciences as HHMI investigators, and who received prestigious early career prizes.

Our results support the hypothesis that appropriately designed incentives stimulate exploration. In particular, we find that the effect of the HHMI program increases as we examine higher quantiles of

the vintage-adjusted distribution of citations. Our preferred econometric estimates imply that the program increases overall publication output by 39 percent relative to early career prize winners; the magnitude jumps to 96 percent when we hone in on the number of publications in the top percentile of the citation distribution. Symmetrically, we also uncover robust evidence that HHMI-supported scientists “flop” more often than the control group: they publish 35 percent more articles that fail to clear the (vintage-adjusted) citation bar of their least well cited pre-appointment work. This suggests that the HHMI investigators are not simply stars on the rise who are anointed by the program, but rather they appear to place more risky scientific bets after their appointment.

These findings are important for at least two reasons. First, they demonstrate the impact of nuanced features of research contracts for the rate and direction of scientific progress. Given the prominent role that scientific change is presumed to play in the process of economic growth, this has important implications for the organization of public and private research institutions. Second, they offer empirical support for the theoretical model developed by Manso,<sup>7</sup> and as such may provide insights relevant to a wider set of industries that rely on creative professionals, ranging from advertising and computer programming to leadership roles at the upper echelons of the corporate world. Many questions remain, and will constitute part of our research agenda going forward.

---

<sup>1</sup> B. Jones, “The Burden of Knowledge and the Death of the Renaissance Man:

*Is Innovation Getting Harder?”* NBER Working Paper No. 11360, May 2005, and *Review of Economic Studies*, 76 (2009), pp. 283–317.

<sup>2</sup> P. Azoulay, J. Graff Zivin, and G. Manso, “NIH Peer Review: Challenges and Avenues for Reform,” NBER Working Paper No. 18116, June 2012, and *Innovation Policy and the Economy*, Volume 13, J. Lerner and S. Stern, eds., forthcoming from the University of Chicago Press.

<sup>3</sup> S. Wuchty, B. Jones, and B. Uzzi, “The Increasing Dominance of Teams in the Production of Knowledge,” *Science*, 316 (May 2007), pp. 1036–9.

<sup>4</sup> P. Azoulay, J. Graff Zivin, and J. Wang “Superstar Extinction,” NBER Working Paper No. 14577, December 2008, and *Quarterly Journal of Economics*, 25 (2010), pp. 549–89.

<sup>5</sup> P. Azoulay, J. Graff Zivin, and B. Sampat, “The Diffusion of Scientific Knowledge across Time and Space: Evidence from Professional Transitions for the Superstars of Medicine,” NBER Working Paper No. 16683, January 2011, and *The Rate and Direction of Inventive Activity: A New Agenda*, J. Lerner and S. Stern, eds., University of Chicago Press, April 2012.

<sup>6</sup> P. Azoulay, J. Graff Zivin, and G. Manso, “Incentives and Creativity: Evidence from the Academic Life Sciences” NBER Working Paper No. 15466, October 2009, and *The RAND Journal of Economics*, 42 (2011), pp. 527–54.

<sup>7</sup> G. Manso, “Motivating Innovation,” *Journal of Finance*, 66 (2011), pp. 1823–60.