NATIONAL CENTER for ANALYSIS of LONGITUDINAL DATA in EDUCATION RESEARCH

TRACKING EVERY STUDENT'S LEARNING EVERY YEAR

A program of research by the American Institutes for Research with Duke University, Northwestern University, Stanford University, University of Missouri-Columbia, University of Texas at Dallas, and University of Washington

AIR

Identifying Teacher Salary Spiking and Assessing the Impact of Pensionable Compensation Reforms in Illinois

Dan Goldhaber
Cyrus Grout
Kristian Holden

WORKING PAPER No. 238-0620-2 • April 2021

Identifying Teacher Salary Spiking and Assessing the Impact of Pensionable Compensation Reforms in Illinois

Dan Goldhaber
American Institutes for Research/CALDER
University of Washington/CEDR

Cyrus Grout
University of Washington/CEDR

Kristian Holden
American Institutes for Research/CALDER

Contents

Acknowledgments

We are grateful to Robert Costrell, James Cowan, Lisa Dawn-Fisher and Mike Podgursky, for comments. Note that the views expressed are those of the authors and do not necessarily reflect the views of the institutions to which the authors are affiliated, or the state agencies whose data were used in this analysis.

This research was supported by Arnold Ventures through a research grant, *Retirement Incentives and Salary Spiking: Learning from Variation in Pension Rules*. All opinions expressed in this paper are those of the authors and do not necessarily reflect the views of Arnold Ventures.

This research was also supported by the National Center for the Analysis of Longitudinal Data in Education Research (CALDER), which is funded by a consortium of foundations. For more information about CALDER funders, see www.caldercenter.org/about-calder. All opinions expressed in this paper are those of the authors and do not necessarily reflect the views of our funders or the institutions to which the authors are affiliated.

CALDER working papers have not undergone final formal review and should be cited as working papers. They are intended to encourage discussion and suggestions for revision before final publication. Any opinions, findings, and conclusions expressed in these papers are those of the authors and do not necessarily reflect the views of our funders.

CALDER • American Institutes for Research
1400 Crystal Drive 10th Floor, Arlington, VA 22202
202-403-5796 • www.caldercenter.org

Identifying Teacher Salary Spiking and Assessing the Impact of Pensionable Compensation Reforms in Illinois

Dan Goldhaber, Cyrus Grout, Kristian Holden
CALDER Working Paper No. 238-0620-2
April 2021

Abstract

Defined benefit (DB) pension plans incentivize “salary spiking,” where sharp increases in pay are leveraged into significantly higher levels of retirement compensation. While egregious instances of salary spiking occasionally make headlines, the prevalence of salary spiking is poorly understood. Moreover, there is little guidance on the definition of salary spiking behavior and how to identify it. This paper develops an empirical method to quantify the prevalence of salary spiking by identifying cases where end-of-career compensation deviates from the expected level of compensation. We apply this method to teacher pension systems in Illinois to assess the prevalence of salary spiking before and after the implementation of a reform designed to dissuade salary spiking.

1. Introduction

Defined benefit (DB) pension systems tend to incentivize "salary spiking," where sharp increases in pay, typically at the end of a career, are leveraged into significantly higher levels of retirement compensation. As discussed below, there are several reasons to be concerned about salary spiking, but the prevalence and magnitude of salary spiking behavior are poorly understood. The lack of a precise definition of salary spiking behavior and broadly applicable method for identifying it contribute to this situation–a shortcoming we seek to address in this paper.

What we do know is that salary spiking behavior occasionally makes headlines—Mannino and Cooperman (2013) surveyed a collection of financial news articles and investigative reports from the period 2000 to 2012 and identified 57 articles on salary spiking across a range of states and sectors. Many of these articles point to salary and retirement payment figures to highlight the costs of salary spiking. For example, a Miami Beach 911 Call Center Worker with an average annual wage of $60,000 used overtime to work an extra 50 hours a week for their last two years on the job, resulting in an *annual* pension benefit of $150,000 (Smiley & Chang, 2011). Other examples of salary spiking involve large cash-out payments: a Ventura County executive was earning a salary of $228,000 as she approached the end of her career, but she also cashed out roughly $70,000 in unused vacation pay and other benefits (vacation pay, etc.). This $70,000 "spike" (roughly 30% of her salary) increased the value of her *annual* pension by about $50,000 (Saillant, Moore, & Smith, 2014). In other words, the $70,000 spike at the end of her career, in all likelihood, increased the value of her pension by hundreds of thousands of dollars.

These cases highlight several policy concerns associated with salary spiking. One concern is that the types of extreme cases described above, however rare, erode the public's faith in pensions and good governance. It is difficult to consider those examples and believe that the pension systems were originally intended to provide those type of benefits. As noted by California Governor Jerry Brown in his Twelve Point Pension Reform Plan, "That one-year [final-average-salary] rule encourages games and gimmicks in the last year of employment that artificially increase the compensation used to determine pension benefits" (2011, p. 2). In other words, many interpret salary-spiking behavior as an abuse of the system and inherently unfair.[1]

A second policy concern associated with salary spiking behavior matters is that it results in unfunded liabilities, the costs of which are ultimately borne by taxpayers and future employees. Public pension systems in the United States are designed to be funded by employee and employer contributions, paid as a percentage of salary over the course of an employee's career. An end-of-career spike in compensation is likely to result in an increase in the value of an employee's pension that far exceeds any corresponding increase in contributions to the pension system made on the employee's behalf. This creates an unfunded liability. At the state level, the financial liabilities associated with salary spiking can be substantial, contributing to the poor financial condition of many state pension systems.[2] Fitzpatrick (2017) estimated that salary spiking among members of the Teacher Retirement System of Illinois cost the pension fund $115 million per year.[3]

[1] Since the financial liabilities associated with spiking are shared across all members of a pension plan, salary spiking will tend to redistribute wealth from non-spikers to spikers.
[2] Estimates peg total U.S. pension shortfalls at several trillion dollars. See Novy-Marx & Rauh (2011) and Biggs (2015) regarding the magnitude of unfunded liabilities. See Zeehandelaar & Winkler (2013) and Malanga & Mcgee (2018) regarding the financial implications of these large burdens on taxes and spending. See Novy-Marx & Rauh, (2009), Anzia and Moe (2017), and Brown, Clark, & Rauh, (2011) regarding the causes of underfunding.
[3] A growing body of research studies the consequences of unfunded liabilities. For instance, Backes et al. (2016) highlight how a large share of retirement contributions of current teachers are used to pay down

A third policy concern is that the financial incentive to engage in salary spiking behavior may also encourage undesirable patterns of employee behavior. There is little reason to think that either employees or employers benefit from employees concentrating their effort into their final year (or years) of employment or hoarding sick leave for cash out that enables a spike. Instead, this dynamic likely creates externalities across agencies within the public pension system to provide additional benefits while only realizing part of the cost (Fitzpatrick, 2017).

In light of these three concerns, it is not surprising that a wide variety of policy initiatives have attempted to discourage salary spiking. For instance, legislators in Massachusetts passed anti-spiking provisions in Chapter 32 that reduces pensionable compensation when the rate of regular compensation exceeds 100% between two years. The California Supreme Court passed a law in 2013 limiting the types of pay that could be used to boost end-of-career compensation, such as termination pay.[4] The Kentucky General Assembly passed Senate Bill 104 in 2017, which caps year-over-year increases in pensionable compensation at 10%.[5] Lastly, Illinois, the subject of focal state in this study, passed legislation in 1979 to explicitly discourage salary spiking by capping year-over-year increases in pensionable compensation at 20%, and passed additional legislation in 2005 to internalize the cost of salary spiking to school districts for increases over 6%.

To our knowledge, there is no general empirical approach to *define and identify* salary spiking. In fact, there is only one empirical analysis of salary spiking in the literature. Fitzpatrick's (2017) analysis focuses on estimating the causal impacts of the Illinois anti-spiking

unfunded liabilities caused by earlier cohorts. There is also some evidence that unfunded liabilities crowd out expenditures on current salaries (Kim et al., 2020).

[4] See https://www.cacities.org/Resources-Documents/Member-Engagement/Professional-Departments/City-Attorneys/Library/2020/Virtual-Conference/10-2020-Virtual;-Ross-Holtzman-California-Supreme.aspx

[5] https://kyret.ky.gov/Publications/Articles/Pages/Pension-Spiking.aspx

policies on salary patterns of teachers as opposed to identifying cases of salary spiking. Indeed, Fitzpatrick's findings highlight how employees and employers that are likely to sidestep policies that rely on simple "percent growth" definitions of salary spiking. [6]

As we demonstrate using simulations, identifying salary spiking behavior is empirically challenging. Standard, empirical approaches are prone to misidentify individuals as salary spiking when they have higher-than-average salary levels, or greater-than-average salary trends, even when using within-individual variation (e.g., fixed effects). In this paper, we develop a definition of salary spiking and a method for identifying instances of salary spiking that addresses heterogeneity in salary levels and trends and performs well in simulations. This approach is highly relevant for policy because it can be applied to any pension system where detailed compensation records are available.

Applying our method of salary spiking identification to teacher compensation data from two Illinois teacher pension systems – the Teacher Retirement System of Illinois (TRSIL) and the Chicago Teachers' Pension Fund (CTPF) – we measure the prevalence and magnitude of salary spiking and assess the financial implications of spiking.[7] We find that around 40 to 50% of employees have unexpectedly large increases in compensation during their FAS periods. This compares similarly to concurrent work by Shuls & Lux (2019), which applies the empirical method described in our working paper (Goldhaber et al., 2018) to Missouri.[8] Back-of-the-

---

[6] Fitzpatrick shows that many school districts awarded retirement bonuses as large as 20% of salary, resulting in additional liabilities estimated at $116 million per year. When in 2005 the state began billing school districts for any pension liabilities associated with salary growth in excess of 6%, effectively internalizing the cost of salary spiking behavior to the employer, school districts changed their behavior, suggesting that they had been awarding compensation inefficiently.

[7] In contrast, Fitzpatrick (2017) focuses on estimating the causal impact of a salary spiking *policy* and indirectly examines the change in spiking behavior as shown by changes in salary.

[8] Shuls & Lux also raise questions about the influence of "macroeconomic fluctuations" on the prevalence of salary spiking, which we describe and explore in Section 3, Salary Spiking Mechanisms.

envelope calculations imply a cost of about $450 million per cohort of exiting members of TRSIL.

The paper proceeds as follows. Section 1 provides describes the key features of TRSIL and CTPF. Section 2 discusses how and why salary spiking occurs, and Section 3 describes the data. Section 4 presents our empirical approach to identifying salary-spiking behavior, and Section 5 presents simulation results. Section 6 presents our findings and Section 7 concludes.

2. Background

In this section, we describe the key features of two pension systems: the Teacher Retirement System of Illinois (TRSIL) and the Chicago Teachers' Pension Fund (CTPF).

Teacher Retirement System of Illinois. Established in 1939, TRSIL operates pension plans that cover school public school employees throughout the state of Illinois (with the exception of Chicago Public School employees). Most TRSIL members are teachers (about 80%), but because the system also includes school administrators and staff; we refer to members as "employees". We focus on employees enrolled in TRSIL's Tier 1 plan, which includes all employees enrolled prior to 2011.[9] Employees currently contribute 9% of their compensation to the pension fund and employers contribute 0.58%.[10] As of 2018 TRSIL was 40% funded, facing an unfunded liability of $77.9 billion – roughly $486,000 per active member.[11]

A typical DB pension plan pays an annuity in retirement that is a function of a benefit factor, a member's years of service (YOS) and final average salary (FAS):

$$Annuity = B * YOS * FAS \tag{1}$$

[9] TRSIL subsequently introduced Tier 2 which enrolled members hired after 2011. We exclude Tier 2 members from our analysis because our data only cover the period 1992-2012.
[10] Interestingly, in Illinois, the state pays the vast majority of the employer contribution. For more on state subsidies of school district pension costs, See Costrell et al. (forthcoming). Additionally, TRSIL members are not covered by Social Security.
[11] For more details, the 2018 GASB report is available at https://www.trsil.org/sites/default/files/documents/GASB2018.pdf (accessed December 3, 2019).

TRSIL has a more complicated relationship between B, YOS, and FAS. Prior to 1998, the benefit factor depended on an employee's YOS so that the FAS replacement rate, B*YOS, is equal to:

- 1.67% for each of the first 10 years
- 1.9% for each of the second 10 years
- 2.1% for each of the third 10 years
- 2.3% for each year over 30 years

For example, an employee with 30 years of service by 1990 would have an annuity of (where $0.567 = 0.0167 * 10 + 0.019 * 10 + 0.021 * 10$). Note that, because the benefit factor increases with years of service, TRSIL plans have even more "backloading" of retirement compensation towards employees who serve long careers relative to plans in other states with constant benefit factors (Costrell and Podgursky 2009). After 1998, TRSIL adopted a fixed benefit factor of 0.022 for all newly accrued YOS, and members were given the opportunity to upgrade the benefit factor applied to previously accrued YOS to 0.022 (see Fitzpatrick (2015) for a detailed discussion). Annuity values are capped at 75% of FAS, at which point, increasing pension annuities is only possible by increasing FAS.

FAS is equal to an employee's pensionable earnings during his or her four highest consecutive years of compensation.[12] Employees can retire with full benefits at age 62 with 5 or more YOS, age 60 with 10 or more YOS, or age 55 with 35 or more YOS. Tier 1 members can also retire early at age 55 with 20 or more YOS with a reduced benefit. Under DB pension systems rules, it is generally the case that not all types of compensation are "pensionable" (*i.e.*,

[12] If these contain a partial year of work, TRSIL adjusts FAS upward as if the employee worked a full year. This may cause reported compensation to understate pensionable earnings, but that said, the average FTE in the analytic sample is quite high at 0.98 so this likely has little practical impact.

eligible for inclusion in the calculation of FAS). A pension plan's definition of pensionable compensation naturally has implications for how employees and employers might engage in salary-spiking behavior. Under TRSIL, the definition of pensionable compensation is relatively expansive, including for instance, salary earned for performing extra duties, bonuses, retirement incentives, severance payments, and payments for unused vacation or sick leave paid or due and payable along with or prior to your final paycheck for regular earnings. Compensation such as workers compensation, payments made after retirement or for work done during retirement, are not pensionable.[13]

Illinois introduced two policies that attempt to limit salary-spiking behavior. In 1979, the state placed limits on pensionable compensation when for year-over-year increases in compensation that are greater than 20%, which we call the "20% growth cap".[14] For example, when an employee has 25% salary growth, the state only allows pensionable compensation to be calculated using 20%. In 2005, the state introduced a new policy that introduced "excess compensation" billings when employees have increases in compensation over 6%, while allowing these increases to be pensionable. The excess compensation rule created a clear incentive for school districts to avoid compensation structures that would result in employees experiencing salary growth in excess of 6% at the ends of their careers, whether through the payment of bonuses or compensation for unused leave. As documented by Fitzpatrick (2017), the

[13] A complete list of pensionable and non-pensionable compensation is available pages 5-6 of the Tier 1 Member Guide available at https://www.trsil.org/sites/default/files/documents/Tier1MemberGuide-August2019-web_2.pdf (accessed December 3, 2019).

[14] The 20% Growth Cap affects pensionable compensation, and likely should be assessed when reporting financial liabilities. In practice, this appears to make little difference overall as relatively few people have gains much above 20 percent. We have considered spiking magnitudes that adjust for 20 percent gains by replacing any value with 1.2 times compensation in the prior year and find that the average magnitude changes by about $700 on a base of around $10,000.

2005 rule changed prompted many school districts to spread smaller retirement bonuses across multiple years of service in order to avoid triggering excess compensation billings.

Chicago Teachers' Pension Fund. Chicago Public Schools (CPS) operates a pension system called the Chicago Teachers' Pension Fund (CTPF) that is separate, though similar, to TRSIL. The Tier 1 plan considered in this study has the same annuity formula (including the benefit formula and upgrade option described above), FAS period, and benefit cap at 75% of FAS. The retirement eligibility rules are slightly different, with CTPF requiring 20 YOS at age 60 instead of 10 YOS for full retirement eligibility. As of June 2018, the CTPF was 50.1% funded with $10.9 billion in unfunded liabilities – roughly $377,000 per active member. Employee contributions to the system are equal to 9% of compensation.

One important difference between TRSIL and CTPF is that CTPF covers only one employer, CPS. Any additional pension income earned from salary spiking in CPS contributes directly to the liabilities of CTPF. Thus, CPS as an *employer* is fully incentivized to prevent salary spiking, while CPS *employees* face similar incentives to salary spike. This contrasts with *other employers* in Illinois*,* who can pass the burden of additional unexpected benefits to the state. It is not surprising then, that CTPF does not have a similar excess compensation rule as TRSIL to internalize the costs of salary spiking across employers. In the section below, we discuss these unique incentives by defining different types of salary spiking and how this might occur in Illinois.

Another important difference between TRSIL and CTPF is the definition of pensionable compensation. As described above, TRSIL has a relatively expansive definition with retirement/severance bonuses explicitly pensionable while in contrast, CTPF has a more restrictive definition. For CTPF, compensation such as merit, longevity, and retention bonuses are pensionable compensation, but explicit retirement, severance, and lump-sum payouts are not pensionable.[15] The more restrictive definition of pay for CTPF may reflect the alignment of incentives for managing pension liability between Chicago Public Schools and CTPF. In particular, a review of CBAs for non-CPS school districts indicates that excess compensation policies likely motivated districts to redefine compensation as not pensionable. Thus, one may view this tightening of definitions as part of the alignment of incentives between employers and pension agencies.

3. Salary Spiking Mechanisms

Before considering what patterns of compensation constitute salary spiking, it is worth thinking carefully about both *why* and *how* salary spiking might occur. We discuss these two points in turn.

As described above, salary spiking is the act of increasing an employee's pensionable compensation with the intent of boosting pension income, and the incentive to spike arises from the fact that a one-time increase in compensation can be leveraged into a higher level of pension income for the duration of one's retirement. For example, an employee in a plan with a four-year FAS averaging period and a 60% replacement ratio (e.g., $annuity = 0.02 * 30\ YOS * FAS$) who receives a one-time spike in pensionable compensation of \$20,000 will increase his FAS by \$5,000 and receive an additional \$3,000 in each year of retirement. Assuming a 4% discount rate, this would translate into additional pension income worth about \$40,000 in present value terms.[16]

---

[15] The definition of pensionable compensation in CTPF is compensation paid for service completed during normal school hours. See https://www.ctpf.org/sites/main/files/file-attachments/admin_rules_-_salary.pdf, retrieved 1/7/2020 for more details.
[16] To streamline discussion, we do not address other costs associated with awarding an employee additional compensation (e.g., pension contribution, taxes, or overhead) since they must be paid whether or not the additional compensation is paid during an employee's FAS averaging period.

This incentive to spike is likely to influence the behavior of both employees and their employers by encouraging them to concentrate compensation into the employee's FAS period. From the perspective of the employee, the incentive to spike compensation is fairly obvious. Continuing with the example above, the value of receiving an additional $20,000 of compensation outside of the FAS period is simply $20,000. Within the FAS period, however, its value is equal to $20,000 in current compensation plus additional retirement income worth $40,000.

To think about an employee's incentive to salary spike more abstractly, consider **Figure 1**. It represents a hypothetical budget constraint for an individual choosing between labor and leisure in her final year of work under a traditional DB pension plan (Panel A), and a plan without a FAS period (Panel B). The budget constraint in Panel A is kinked at the point where an employees' potential current compensation is higher than at least one of the years that would otherwise be included in her FAS calculation. To the left of this point, the employee's rate of compensation is higher because increased effort affects both current compensation and retirement compensation. Under a higher rate of compensation, employees with stronger preferences for consumption over leisure and effort will tend to increase their supply of labor inside the FAS period, resulting in a spike in compensation. This contrasts with plans that do not calculate FAS, as shown in Panel B. Under a DC plan, if labor/leisure preferences are smooth over time, one would not expect to observe a discontinuous shift in effort at the end of an employee's career because the marginal rate of substitution between leisure and effort is equal to the (constant) rate of compensation. In contrast, one can show that under a DB pension system, employees cannot maximize utility at the same level of leisure and effort and will maximize utility by either *decreasing or increasing compensation*.[17]

Figure 1. Illustrations of employee salary spiking incentives

Panel A: Kinked budget constraint under DB pension FAS rules

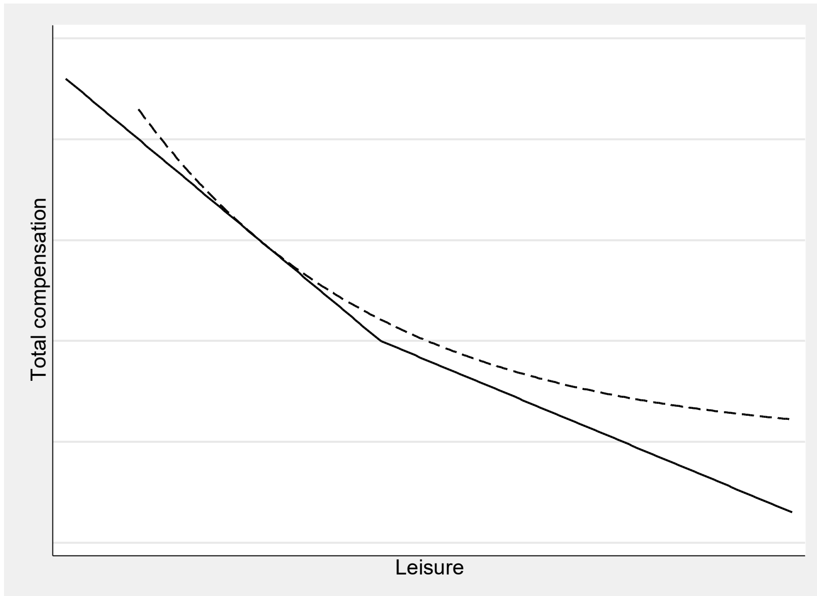

Panel B: Linear budget constraint without pension FAS rules

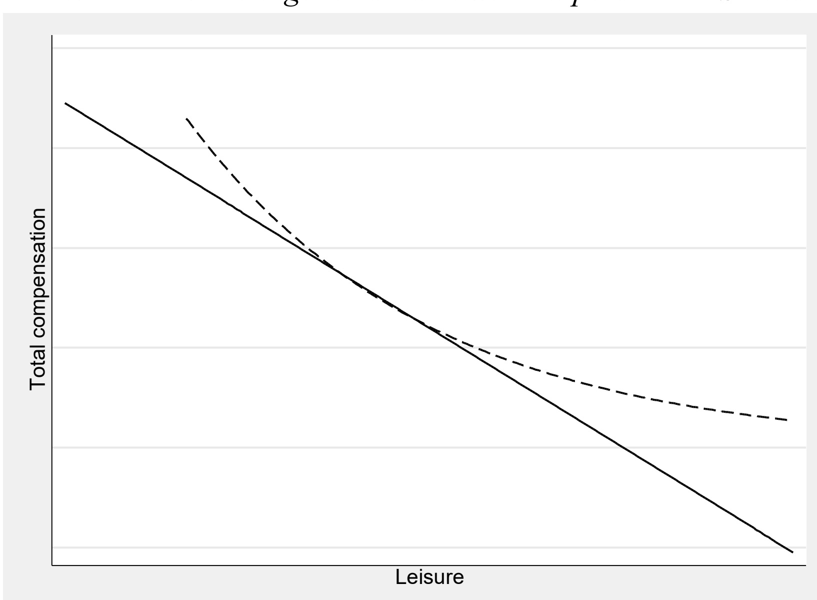

Notes: Panels represent a hypothetical budget constraint for an individual choosing between labor and leisure in their final year of work. In Panel A, the kink is at the point

[17] Note that an optimal allocation of labor and leisure in **Figure 1** requires the marginal rate of substitution between compensation and leisure to be equal to the slope of the budget constraint, but at the kink point, the budget constraint has two slopes, so no single marginal rate of substitution will satisfy both.

at which FAS begins to increase as salary in the final year exceeds one of the prior FAS years. Panel B represents a system without an FAS period.

From the perspective of the employer, the incentive to salary spike is less obvious since the employer does not directly benefit from the influence of a salary spike on an employee's pension income. However, returning to the example above, consider that an employer can effectively award $60,000 in additional compensation while only expending only $20,000 if that $20,000 awarded inside an employee's FAS period. If the employer is one of many employers in the pension system, the cost of funding the $40,000 in additional pension liabilities will spread out across enough actors that the employer will bear very little of it.

Given that both employees and employers have an incentive to engage in salary-spiking behavior, one might expect the *how* of salary spiking to involve both parties. As such, we characterize salary spiking related to retirement bonuses as being both *institutional spiking* – characterized by actions taken at an institutional level (e.g., a school district or state agency) – and *employee spiking* – characterized by actions taken at the level of individual employees.[18]

For *institutional spiking,* a natural institution to consider is within the school district as compensation for employees is generally determined by negotiation with school districts (Strunk et al. 2018). In this case, employees could push for compensation via additional duties and responsibilities, by bargaining for changes to the salary schedule that would affect all employees, or for compensation like retirement or longevity bonuses. The use of retirement bonuses in TRSIL are an excellent illustration of this. As documented by Fitzpatrick (2017), rules guiding the provision of retirement bonuses are stipulated in collective bargaining agreements (CBAs)

[18] Alternatively, employers may not be as distinct from employees in public education. For instance, research by Moe (2006) suggests that district employees have strong incentives to get involved in school-board politics to elect candidates aligned with their own interests.

negotiated between school districts and teachers' unions. Other institutions may exist, as employees may influence state-level policy, such as how Illinois state policy recently dropped a proposed policy to charge excess compensation for compensation growth over 3% after concerns were expressed from teacher organizations.[19]

Given the constraints set at the *institutional level*, employees decide whether to seek out additional compensation that will increase pensionable compensation. *Employee spiking,* for instance, can occur when employees choose whether to pursue additional duties and responsibilities to increase their pensionable compensation. In the example of retirement bonuses in TRSIL, not every employee within a district that provides retirement bonuses will ultimately receive one.

Finally, some salary spiking may occur as *incidental spiking*. For instance, suppose that a state increases compensation for all employees in a given year. Those who coincidentally retire in the same year will all benefit from a spike in salary, but they did not *intend* for this effect; we define this unintended outcome as *incidental salary spiking*. This contrasts with the individual and institutional level mechanisms discussed above which speak to the intent. To be clear, the data-driven methods for identifying salary-spiking behavior that we advance below do not distinguish between different types of salary spiking. Rather, our approach looks for deviations in an employee's historical pattern of compensation and is agnostic to the motivations and mechanisms underlying any such deviation. While this is a limitation from a *behavioral* perspective, it is an advantage from an actuarial standpoint because spiking can have important impacts on unfunded liabilities regardless of their intent.

[19] See https://www.forbes.com/sites/ebauer/2019/06/04/hidden-in-the-legislate-a-thon-illinois-restores-pension-spiking/#5447acd92193, retrieved 1/14/2020.

4. Data

Employment records were obtained from the Illinois Teacher Service Record (TSR) via a Freedom of Information Act (FOIA) request from the Illinois State Board of Education (ISBE). These records provide annual employment data for teachers, administrators, and other school employees from 1991-1992 to 2011-2012 including information on annual compensation and experience. According to ISBE, compensation and experience values are representative of pensionable compensation and service credit (e.g., experience that affects pension benefits as reported in Equation 1). [20]

As described below, our empirical approach requires observing patterns of compensation over a sufficiently long period of employment to be able to establish a pattern of compensation and identify any deviation from that pattern of compensation. Therefore, we focus on individuals with at least 10 observations in our data.[21] We also limit the sample to those who are observed exiting employment. Overall, these restrictions define a sample of 59,724 TRSIL and CTPF employees, with about 5,400 employees exiting each year between 2000-2001 and 2010-2011. Of the 271,560 unused employees, about half are not observed exiting the sample by 2012 (132,559) and should clearly be excluded as they are continuing work, while the rest are not observed in for at least 10 years of data. This latter group includes employees who appear to exit the sample with less than 10 YOS (about 57%) and censored records of employees with higher levels of experience who exit within 10 years of the start of our data in 1992.

[20] Fitzpatrick (2017) states that TSR changed salary reporting practices for in 2003 by requiring districts to report salary earned over the summer. To the best of our knowledge, we are not able to verify this, as our communications with ISBE indicate that summer earnings are included in all TSR years, and they are not aware of any change in 2003.
[21] As described below, 10 observations in our most restrictive model allows for only 3 degrees of freedom and individuals with fewer observations would likely have very poorly fit data.

A limitation of the data is that it does not contain a variable that uniquely identifies unique employees. Therefore, we link individual records across years using first, middle, and last name, and the institution where they received their baccalaureate degree. We adopted a conservative approach, keeping only exact matches, with one caveat.[22] The year-over-year match rate ranged between 90 and 95% which is consistent with the 8% rate of teachers leaving the profession as reported by NCES for 2012-2013.[23]

Descriptive statistics for employees included in the analytical sample are presented in **Table 1**. The first column shows means for all employees, and the second and third columns report statistics for employees in CTPF and TRSIL, respectively, with the last column reporting the difference between TRSIL and CTPF. All statistics represent employees' characteristics as of their final year of employment. Consistent with the presence of salary-spiking behavior, we see that average compensation growth is higher during employees' final two years of service (T and T-1). The increase in growth rates is greater in non-CPS districts.

---

[22] Between 1998-99 and 1999-2000, the structure of name fields changed from having one field containing first, middle and last name to three separate fields. Because name order is not preserved consistently across employers, the exact match rate is less than 40 percent. For this period, we use a fuzzy match algorithm (85% of records can be matched with a score greater than 0.98) and to link records for any individual with records before and after 1999. Records are matched using the user-written Stata command reclink; "Or-blocking" is used so that only records with matching names or universities are considered.

[23] https://nces.ed.gov/programs/coe/pdf/coe_slc.pdf

Table 1. Employee characteristics as of last year of employment

	All	CTPF	TRSIL	TRSIL – CTPF*
Year-over-year salary growth				
T (Final year)	7.2%	2.4%	8.5%	6.1%
T-1	7.8%	5.8%	8.4%	2.6%
T-2	5.6%	5.0%	5.8%	0.8%
T-3	5.5%	6.1%	5.4%	-0.8%
Demographics				
Female	0.71	0.75	0.70	-0.05
White	0.83	0.44	0.95	0.51
Employment characteristics				
Advanced degree	0.72	0.68	0.72	0.05
Salary	78,178	75,704	78,878	3,174
	(32,507)	(25,093)	(34,281)	
Years of service (YOS)	28.07	27.78	28.15	0.37
	(8.25)	(9.34)	(7.91)	
Job position				
Administrator	0.08	0.09	0.08	-0.02
Teacher	0.78	0.77	0.79	0.02
Other position	0.14	0.14	0.13	0.00
Observations	59,724	13,163	46,561	

*Notes: Samples come from a panel of employment data from 1991-1992 to 2011-2012, limited to employees with at least 10 years of observable consecutive employment at the end of their careers (see sections 4 and 5 for further discussion). CTPF is the Chicago Teachers Pension Fund and includes employees who work for Chicago Public Schools at some point in their observed career. TRSIL is the Teacher Retirement System of Illinois, which covers all other employers in the state, and includes employees who are never observed in CPS. *All differences between CTPF and TRSIL are statistically significant at the 0.01% level.*

Apart from having very different racial demographics, employees in CPS and CTPF have similar characteristics. The majority are female, tend to hold an advanced degree (Masters or Doctorate), earn about $78,000 in nominal dollars, and exit with about 28 years of service.[24] The sample is composed primarily of teachers. Relatively few are administrators (e.g. principals,

[24] As expected, there is a significantly higher amount of experience in this sample because we focus on exiting employees as opposed to the average employee in TRSIL and CTPF.

assistant principals, and support staff) or employees serving in other roles (e.g., counselors, janitorial).

5. Empirical Approach

On a case-by-case basis, distinguishing between a compensation increase that constitutes salary spiking and one that does not, may not be terribly difficult. However, it would be extremely resource intensive to gauge the system-wide prevalence and magnitude of salary-spiking behavior in this manner. As noted above, the empirical approach we advance below to identify salary-spiking behavior is agnostic as to *why* or *how* an employee experienced any particular change in compensation. In this sense, we do not precisely identify who is, or is not, spiking. The advantage of our approach is that it establishes a definition of salary spiking that can be consistently applied to any pension system using readily available administrative data.

We propose the following empirical definition of salary spiking: an employee is salary spiking if his end-of-career compensation significantly and positively deviates from his prior pattern of compensation. To put this definition into operation, we define a range of end-of-career compensation that falls "within expectations" given a employee's preceding levels of compensation and comparing that employee's *actual* end-of-career compensation to the range of *expected* compensation.[25] To define a range of end-of-career compensation that is within expectations, we use simple forecasting methods. Specifically, we regress compensation on years to exit:

(2) $$C_{it} = \alpha_i + \beta_i' f(Year_{it}) + \varepsilon_{it}, \;\; t < T,$$

[25] We focus on end-of-career salaries, and the final four years of employment in particular, because roughly 90 percent of TRSIL employees earn their highest four years of compensation during their final four years of employment. Hence, for the great majority of employees, the FAS period is the final four years.

where C_{it} is employee i's compensation in year t, $Year_{it}$ is the school year, $f(\cdot)$ is a polynomial function of $Year_{it}$, and β_i' is a vector of coefficients estimated separately for employee i. Note that the forecast model estimates separate intercept and slope coefficients for each employee. In other words, the model is fully interacted such that it is equivalent to estimating a separate regression model for each employee.[26]

The estimated parameters from this model are used to forecast the final years of compensation for employee i. In our preferred specification, we forecast the final four years of compensation (equal to the length of the FAS period):

(3) $$\hat{C}_{it} = \hat{\alpha}_i + \hat{\beta}_i' f(Year_{it}),$$

and the 95% confidence interval for the forecast of $\hat{C}_{iT}$:

(4) $$CI_{it} = \left[\hat{C}_{it}\ ,\ \hat{C}_{it} + T_{0.05} * S_{ie} * \sqrt{1 + h_{it}}\,\right]\ ,$$

where $T_{0.05}$ is the t-statistic for a one-tailed test, S_{ie} is the standard error of the regression for employee i, h_{iT} is the Tth diagonal element of the projection matrix given by $x_t{}'(X'X)^{-1}x_t$, X is the matrix of independent variables in **equation (8)**, and x_t is the tth row of X.[27] For our primary specification, we adopt a quadratic functional form and use a 95% confidence level in defining the T-statistic.[28]

Given the range of expected compensation defined by CI_{it}, we define the indicator variable $Spike_i$ as follows:

---

[26] It may be tempting to estimate a pooled regression with additional controls for employee and employer characteristics in the interest of improving precision but as discussed below, doing so would advance a conceptual definition of salary spiking that differs from the definition proposed above.

[27] See Greene (2003) for a detailed discussion of this statistic on page 111.

[28] Regarding the functional form of the regression models, our primary specification is a simple quadratic polynomial in school year. This specification has been widely used to fit age-earnings profiles of workers (e.g., Mincer 1974), though we are sensitive to concerns that such models may not provide an appropriate fit (Murphy and Welch), and so, we present a figure in Appendix A indicating that functional form is likely not an issue relative to the magnitudes of salary spiking we find.

(5) $$Spike_{it} = 1 \textbf{ IF } C_{it} > CI_{it}\ ; \textbf{ ELSE } Spike_{it} = 0$$

In other words, an end-of-career increase in compensation is characterized as salary spiking when the difference between actual and forecast compensation is positive and statistically significant. Note that we choose a 95% level of confidence out of convention, which places a high degree of certainty on whether an employee's compensation deviates from prior patterns of compensation. For example, a 95% confidence level implies that the average increase in compensation in a final year of employment needs to be larger than $8,227 to be considered salary spiking, while a less conservative 80% confidence level implies an average of $6,117. We report qualitatively similar results using this lower confidence level in Appendix A.

To illustrate our approach, we apply our method to two hypothetical employees in **Figure 2**: an employee who is (by our definition) salary spiking and an employee who is not. For each individual, we forecast a range of expected salaries in their final year of service (2009-2010) defined by **equation (10).** The magnitude of the salary spike is defined as the vertical distance between actual compensation and forecast compensation. In Panel B, the hypothetical employee's final compensation in 2009-2010 falls within the expected range of compensation, and he is not identified as salary spiking. While the predicted compensation is higher than actual compensation, we define the magnitude of the salary spike as zero because they are within the CI.

Figure 2. Example of actual and forecasted salary for representative employees

Panel A: Example with salary spiking

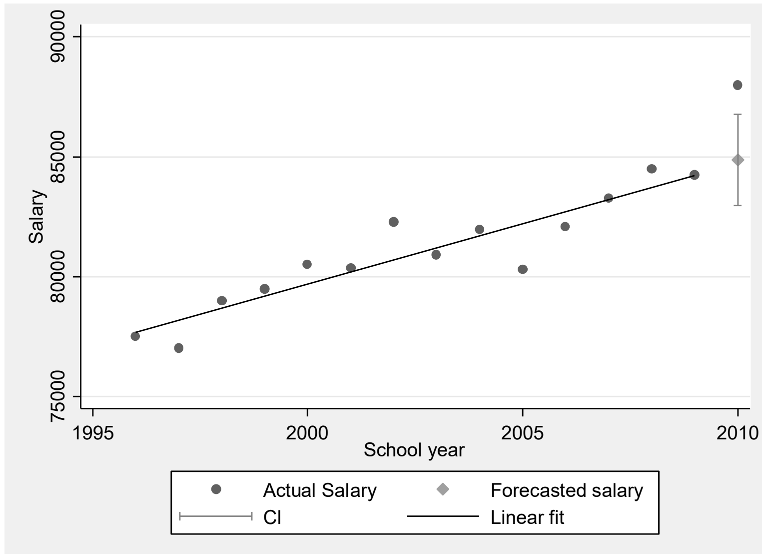

Panel B: Example without salary spiking

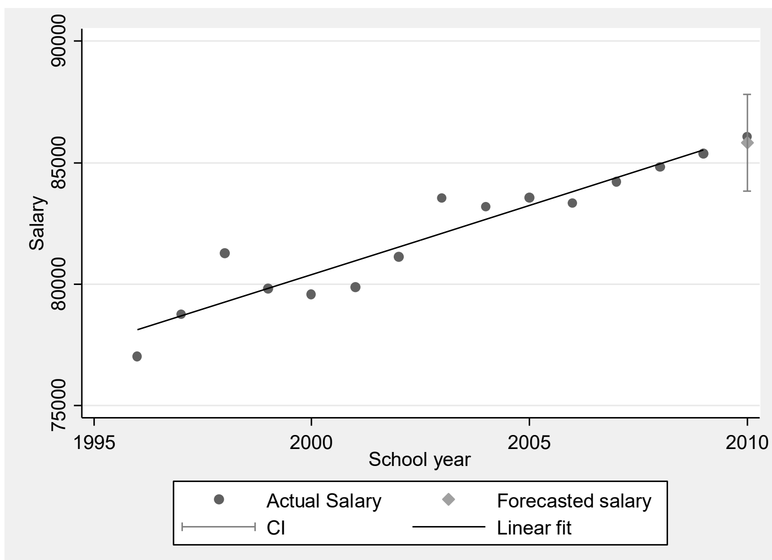

Notes: The figures represent three hypothetical individual's service between 1996 and 2010: Panel A, spiking, Panel B, not spiking, Panel 3, ambiguous spiking. The line indicates fitted values from a regression model excluding 2013, and the diamond point in 2013 indicates the "expected" salary forecast. The confidence interval is constructed from Equation (3) using the standard error of the forecast. The final data point in 2013 lies outside of this range indicating that the individual has a positive salary spike.

In specifying **equation (8)** above, we must make a number of practical considerations related to the treatment of part-time employment, the number of years to include when estimating **equation (8)**, and the range of years over which to forecast compensation. First, regarding panel length, our primary model specification uses all available years of employment data for each employee—either to form predicted compensation or as a comparison to a prediction. This approach has the appeal of using the full set of information available. That said, researchers have argued that there is a bias-variance tradeoff in forecasting if more recent years of data contain better information for predicting final compensation.[29] As such, we also estimate models restricted to each employee's 10 most recent years of service.

Regarding the number of years to forecast, the FAS averaging period is four years for both TRSIL and CTPF and there is an incentive to spike salary during each of these years. Therefore, we forecast expected compensation for an employee's final four years of service because including FAS years in the regression model may bias $\hat{\alpha}_i$ and $\hat{\beta}_i$ if an employee is spiking in years T-1, T-2 or T-3.[30] While an employee's FAS period may plausibly occur outside of her final four years of service (if end-of-career compensation was declining, for instance), this is rarely the case among members of TRSIL – 98% of members experience their highest earning period during their final four years.

As noted above, our regression approach is equivalent to estimating independent regression model for each employee. In the pursuit of greater precision, it is tempting to pool observations in order to leverage the large size of the analytic sample. However, a pooled model would be conceptually incongruous with how salary spiking occurs. Specifically, it would characterize salary spiking as having a higher than expected end-of-career compensation given one's observable characteristics, whether or not that corresponded with an increase in pay. Furthermore, because we are interested in testing whether an employee's end-of-career

[29] See Clark and McCracken (2009), or Greene (2003, page 112) for examples.
[30] In an Appendix, we present results where these years *are* included in the regression models and fewer years are forecast.

compensation significantly deviations from prior patterns of compensation, applying estimates of variance derived from the overall sample will tend to be biased.

6. Simulations

Estimating salary spiking behavior turns out to be a more challenging empirical task then one might imagine. In particular, as we illustrate below, using typical estimation methods to try to identify salary spiking behavior is likely to lead to biased estimates of spiking. We highlight these challenges by simulating data under two hypotheses: 1) no *systematic* increase in final year salary (which we refer to as "no spiking")[31], and 2) a latent propensity model that induces 40% of the population to spike in their final year.[32] We estimate spiking behavior using three different models—A) ordinary least squares (OLS), B) OLS with individual fixed effects, and C) our preferred specification (described above). We then use each model to calculate predicted values to forecast salary in each individual's final year, construct a forecast interval, and determine whether each individual in the sample exceeds this interval.

The key problem, as we show below, occurs when salary patterns differ across individuals in the sample, particularly when individuals have different salary trends. Said simply, OLS models will incorrectly identify individuals with higher-than-average *trends* in salary as having salary spiking behavior—which can lead to identifying salary spiking behavior where none is present or fail to identify individuals who salary spike. We demonstrate this idea by estimating the three salary spiking models described above on data created by four types of data generating processes (DGPs): 1) a common linear trend and intercept, 2) heterogeneous

[31] By systematic, we mean distinct from random chance where some individuals could have particularly large increases in salary in their final year, even though each year is generated from the same distribution. In particular, five percent of individuals will have an unexpectedly large final salary relative to the rest of the sample.
[32] The latent propensity model is specified so that each individual in the sample has an underlying propensity variable, distributed uniform from zero to one, and for those who fall below 0.4, their salary in the final year is 10 percent higher than their previously generated value.

intercepts but common trends, 3) heterogeneous trends but a common intercept, and 4) heterogeneous trends and intercepts. In addition to showing the proportion of individuals identified as salary spiking, we also show the error rate in Panel B (both non-spiking individuals incorrectly identified as spiking, or spiking individuals who are not identified).

Table 2 presents results from the simulations with each column showing a different DGP, and the two panel show settings without and with salary spiking. Starting with Panel A, column 1 represents a baseline case, where OLS is an appropriate strategy by construction—indeed, both OLS and OLS with individual fixed effects perform well, with relatively few individuals identified as salary spiking in Panel A (about 4%). Columns 3 and 4 have DGPs with heterogeneous trends, and heterogeneous trends and intercepts, respectively. OLS and OLS with individual fixed effects greatly overstate the amount of salary spiking when none is present (around 12 to 16%). Looking more closely, we find that *all* of these falsely identified individuals have above average salary trends. In contrast, the preferred specification has a low false identification rate of about 5 to 6%, and only about half of these (55%), have above average salary trends.

Table 2. Simulation results by DGP, estimation strategy, and nature of spiking

DGP:	(1)	(2)	(3)	(4)
Panel A: Hyp. 1, No spiking in data generating process				
Proportion incorrectly identified as spiking:				
True	0.00	0.00	0.00	0.00
A) OLS	0.04	0.04	0.16	0.12
B) Fixed effects	0.04	0.04	0.16	0.15
C) Preferred specification	0.06	0.06	0.06	0.05
DGP with:				
Het. intercepts	No	Yes	No	Yes
Het. trends	No	No	Yes	Yes
Number of individuals	2000	2000	2000	2000
Panel B: Hyp. 2, 40% of individuals spiking in data generating process				
Proportion identified as spiking:				
True	0.39	0.41	0.41	0.40
A) OLS	0.41	0.08	0.20	0.15
B) Fixed effects	0.42	0.43	0.24	0.22
C) Preferred specification	0.42	0.44	0.43	0.41
Error rates:				
A) OLS	0.02	0.38	0.41	0.39
B) Fixed effects	0.03	0.02	0.37	0.36
C) Preferred specification	0.03	0.03	0.04	0.04
DGP with:				
Het. Intercepts	No	Yes	No	Yes
Het. trends	No	No	Yes	Yes
Number of individuals	2000	2000	2000	2000

Notes: Each column represents a separate Data Generating Process (DGP) that introduce progressively more heterogeneity in salary patterns across individuals. Panel A represents a setting where there are no systematic changes in salary, and Panel B introduces a large salary spike in about 40% of the population using a latent variable approach, and "True" represents the number of individuals who are salary spiking. Three empirical approaches are applied: OLS regression, a teacher-level fixed effects regression, and our preferred empirical approach, described in our Empirical Section above. Error rates are calculated when the predicted value of spiking does not agree with the true spiking propensity as determined by the DGP.

Next, we turn to Panel B where salary spiking occurs for 40% of the population on average. In column 1, we see that slightly more than 40% of the sample are identified as salary spiking in each model (about 41 to 42%). Moreover, the error rates in Panel B are quite low under the baseline—most of the individuals who are identified are in fact spiking. Column 2 represents a DGP where individuals have heterogeneous levels (intercepts), and we find that the OLS model now greatly under-identifies spiking on average (8%), while results are quite similar for OLS with individual fixed effects and our preferred specification. Under columns 3 and 4 (DGPs with heterogeneous trends) we find that both OLS and OLS with fixed effects tend to under-identify spiking on average. Perhaps most importantly, the OLS and OLS with individual fixed effects have very large error rates in these last two columns—about 36 to 41% *of the sample* is misidentified, suggesting that these models do not identify spiking behavior well at all. In contrast, the preferred specification performs very well with an error rate of about 3 to 4 percent.

Expanding on the idea of error rates shown in the previous table, we proceed by showing the type of errors for each empirical strategy by presenting the proportion of observations classified as spiking or not relative to whether the observation is truly spiking, as determined by the DGP. These results are shown in **Table 3**, which includes four possible patterns of observations: 1) not identified as spiking and not actually spiking, 2) not identified as spiking but actually spiking, 3) identified as spiking but not actually spiking, and 4) identified as spiking and actually spiking. Cases 1 and 4 represent correct predictions by each model, while cases 2 and 4 represent misidentification.

Table 3. Simulation results comparing identified versus actual spiking by DGP and estimation strategy when spiking is present

	DGP1		DGP2		DPG3		DGP4	
	No Spike	Spike	No Spike	Spike	No Spike	Spike	No Spike	Spike
OLS								
No Spike	0.59	0	0.57	0.36	0.49	0.31	0.53	0.32
Spike	0.02	0.39	0.03	0.05	0.1	0.1	0.07	0.08
OLS with FE								
No Spike	0.59	0	0.57	0	0.49	0.27	0.51	0.27
Spike	0.03	0.39	0.02	0.41	0.1	0.14	0.09	0.13
Preferred								
No Spike	0.58	0	0.56	0	0.56	0.01	0.57	0.02
Spike	0.03	0.39	0.03	0.41	0.03	0.4	0.03	0.39

Notes: Each column represents a separate Data Generating Process (DGP) that introduce progressively more heterogeneity in salary patterns across individuals, as described in Table 2 Panel B; in other words, the DGP introduces a large salary spike in about 40% of the population using a latent variable approach. Three empirical approaches are applied: OLS regression, a teacher-level fixed effects regression, and our preferred empirical approach, described in our Empirical Section above.

For DGP1, and as presented in the previous table, all models tend to be accurate—for example, OLS predicts 59% of the sample is not spiking correctly, and 39% of the sample as spiking correctly, with only 2% error. Under this simple DGP, all true spiking cases are correctly identified. This is because all individuals in DGP1 have the same salary trend and level, which makes for highly accurate predictions of expected salary in the final year relative to the size of the spike, so that all deviations created by the DGP are highly statistically significant.[33] In contrast, all errors are of type 3; in other words, the empirical models incorrectly identify individuals as spiking when they are not.

DGP2, which introduces heterogeneous intercepts, shifts a large proportion of correctly identified spiking cases to be not identified for the OLS model (only 5% correctly identified as

[33] Choosing alternate DGP parameters, such as a smaller spiking magnitude or introducing more noise, would negate this outcome. This latter case is introduced in DGP2-DGP4 where salary levels and/or trends have much more variation than in DGP1.

spiking and 36% identified as not spiking when they are truly spiking). Said differently, the OLS model does not distinguish different levels of pay across individuals from spiking behavior. As before, OLS with FE can separate levels of pay from spiking behavior under DGP2, and again correctly identifies all cases of true salary spiking.

DGP3 and DGP4 represent cases that are problematic for both OLS and OLS with FE. For OLS, there is a similar pattern to DGP2, where a large proportion of true spiking cases are not identified (31% in DGP3, 32% in DGP4). Moreover, there is an increase in the number of cases that are incorrectly identified as spiking when they are not (10% and 7% in DGP3 and DGP4, respectively). OLS with FE performs slightly better for both DGP3 and DGP4 relative to OLS, but is far below the error rate of the preferred specification.

Overall, these simulation results suggest that the preferred model has far better performance than the other models under DGPs with either heterogeneous intercepts, heterogeneous trends, or both—regardless of whether spiking is present or not. However, under DGP1, the preferred specification has slightly higher error rates than these models, and as such, it's worth asking whether it is reasonable to assume whether *empirical data* tends to resemble DGP1, or whether it resembles DGP2, DGP3, or DGP4. We consider this by estimating a simple F-test comparing the OLS model to one that includes individual-specific slopes and intercepts using data from Illinois. The F-statistic is 17.01, and has a P-value $< .001$, indicating that we can reject the null hypothesis that individuals in Illinois tend to have the same level and trends in salary.

7. Results

Prevalence and Magnitude of Salary Spiking. Here we present evidence on the prevalence and magnitude of salary-spiking behavior using data from employees in TRSIL and CTPF.

Table 4 considers three samples of employees: exiting members of TRSIL (column (1)), exiting members of CTPF (column (2)), and non-exiting members of TRSIL (column (3))—this group has little incentive to salary spike because they are unlikely to be in their FAS period, and allows for some characterization of increases in salary that are not related to pension benefits. For each sample group, we find that average model fit is quite good, with adjusted R-squared statistics averaging over 0.90. For the TRSIL sample, 95% of the individual regressions have an adjusted R-squared statistic above 0.85. (column (3)).

Among TRSIL employees exiting employment 46% are identified as salary spiking in year T, 40% in year T-1, 26% in year T-2, and only 15% in year T-3 using a 95% confidence interval. In Appendix A, we report qualitatively similar results using an 80% confidence interval, with the proportion of employees identified as salary spiking 12-16%age points higher than that reported in **Table 4**. As described above, we define the magnitude of a spike in compensation as the difference between actual and forecast compensation in year t if an employee is identified as salary spiking in that year, and zero otherwise. Among the TRSIL members identified as salary spiking, the average magnitude of the spikes in compensation (across the four years) is $43,500, corresponding to an increase in FAS of $10,875.[34] These results are consistent with both quantitative evidence of salary spiking (e.g. Fitzpatrick, 2017) and qualitative evidence mentioned above.[35]

[34] For an employee with 30 years of service, this would increase her annual benefit by roughly $7,000.

[35] Given the relatively large effects, it is potentially interesting to examine district-level patterns in compensation in more detail. This is somewhat challenging, though we have explored the Illinois Salary Survey in 2010-11 in Appendix A. Most types of contract features are unrelated to salary-spiking propensities, such as the structure of the district salary schedule, sick leave policies, and merit or performance pay. Perhaps not surprisingly, union organization is correlated with salary spiking (e.g. IEA-NEA & IFT-AFT relative to independent employers).

Table 4. Estimates of the prevalence and magnitude of salary spiking

	TRSIL	CTPF	Non-exit TRSIL
Model Fit (Avg. R-squared)	0.93	0.91	0.90
Coefficient on quadratic term (Avg.)	-0.19	69.98	32.66
Proportion Spiking in			
T (Final year)	0.46	0.20	0.21
T-1	0.40	0.26	0.20
T-2	0.26	0.20	0.18
T-3	0.15	0.17	0.13
Magnitude among spikers (Avg.)	$43,500	$37,519	$47,816
Observations	46,561	13,163	31,861

Notes: Column 1 reports results for TRSIL employees who exit, and column 2 reports results for exiting CTPF employees. Column 3 presents results using non-exiting TRSIL individuals who are observed in all years of the data, and a randomly selected pseudo-final year. Model fit is summarized by adjusted R-squared terms. Each regression includes school year and school year squared, and average coefficients on the quadratic term are reported in the second row. Prevalence is a sum over an indicator for whether an individual has a salary exceeding the CI given by equation (3). The total magnitude of salary spiking is measured as the sum over difference $C_{iT} - \hat{C}_{iT}$ *for each year an individual is identified as spiking. Samples come from a panel of employment data from 1991-1992 to 2011-2012, limited to employees with at least 10 years of observable consecutive employment at the end of their careers.*

Next, we consider CTPF members. Between 17% and 26% are identified as salary spiking during the final four years of service. That the prevalence of spiking among CTPF members is substantially lower than among TRSIL members is consistent with the fact that CTPF members are served by a single employer. In this setting, the employer (CPS) bears the full burden of the costs associated with salary spiking – it is not shared among a large number of employers as is the case in TRSIL.

Lastly, we consider TRSIL employees who are *not* identified as exiting employment (i.e., are still employed as of 2011-2012) and who have no pension-driven incentive to spike salary.[36] We find that between 13 and 21% of employees are identified as having unexpectedly large increases in salary in a particular year – similar to the proportion of exiting CTPF members identified as salary spiking. These results suggest that a significant proportion of the salary spiking identified among exiting members of TRSIL and CTPF may be incidental (i.e., driven by fluctuations in compensation that are independent from an influence of the pension system).

Another potential reason for the higher-than-expected level of spiking among the non-exiting group is that the estimated models have either persistent bias or incorrectly estimated precision.[37] While we find little evidence of misspecification, it is widely acknowledged in the forecasting literature that forecast intervals tend to have poor coverage rates; for example, Makridakis, et al. (1987) consider simulations using M-Competition data and find that 95% confidence intervals only contain about 80% of observations.[38] This is potentially because these intervals do not account for parameter uncertainty, model misspecification, or changes in the data generating process. Compared to prior studies, such as Hyndman et al. (2002) who find coverage rates between 71 and 87%, we find higher rates of 79 to 87%. Interestingly, the rates reported in column (3) are fairly close to rates found by Williams and Goodman (1971) who also consider a simple linear regression forecast.

[36] We randomly choose a year for each individual to treat as their "false exit" year. To facilitate comparison, the sample consists of TRSIL members and not CTPF.

[37] We also explore the possibility of bias by considering "negative" spiking among nonexiters, and find quite similar proportions to Table 4; the symmetry of this result suggests that estimates of variance overstate precision. A similar proportion of "negative" spiking occurs among exiting employees as well, and are available on request. Alternatively, not accounting for parameter uncertainty would lead to a downward bias in our estimate of sigma, leading to forecast intervals that are too small.

[38] See also Wallis (1974), Newbold and Granger (1974), Gooijer and Hyndman (2006) for a discussion of coverage.

In considering the results for the non-exiting group and what they suggest about incidental spiking, it is important to note that a large spike in compensation inside an employee's FAS averaging period will generate unfunded liabilities regardless of *why* it has occurred. As such, while incidental spiking may not reflect pension-driven behavior, it is still of interest to those concerned about the funding status of a pension plan. Comparing CTPF and non-exiting TRSIL employees suggests that CTPF has few issues with individual or institutional salary spiking, but a substantial amount of incidental salary spiking.

Adoption of the 2005 Excess Compensation Rule. As noted above, in 2005 Illinois adopted a policy that requires TRSIL employers to pay "excess compensation" fees when employee salaries increase by more than 6% during years that contribute to the calculation of FAS. These district fees are calculated by TRSIL as the actuarial value of the increase in expected retirement wealth caused by increases in compensation in excess of 6% with payment due a employee retires from the workforce. In this section, we present descriptive evidence comparing the prevalence and magnitude of salary spiking among TRSIL employees before and after the implementation of the 2005 rule. One challenge is that the rule only became binding once a school district negotiated a new contract with its employees. Because we do not have information on contract timing, we exclude the years 2005, 2006, and 2007 (Fitzpatrick (2017) indicates that the policy was binding for all school districts as of 2008).

Table 5. Estimates before and after 2005 anti-spiking policy for TRSIL employees

	Pre-2005	Post-2007	Post minus Pre
Model Fit (Avg. R-squared)	0.94	0.93	
Coefficient on quadratic term (Avg.)	-50.34	28.19	
Proportion Spiking in			
T (Final year)	0.53	0.32	-0.21***
T-1	0.45	0.28	-0.16***
T-2	0.29	0.20	-0.10***
T-3	0.17	0.13	-0.04***
Magnitude among spikers (Avg.)	$43,554	$44,443	
Observations	30,110	16,829	

Notes: Each column reports statistics for TRSIL employees who have a FAS year in a given period, Pre-2005 and Post-2007. Notably, 378 individuals have a gap in service that causes them to be observed in both the pre- and post-period during their FAS period. See Table 2 notes for a description of these models.

Table 5 summarizes the previously presented regression results for employees enrolled in TRSIL in the pre- and post-policy periods, and provide *descriptive evidence* about the impact of the policy on salary spiking.[39] Wee see that salary spiking is more common in the pre-policy period, with the highest prevalence of spiking in employees' final two years of service. About 53 and 45% of individuals are identified as salary spiking in years T and T-1, respectively. The prevalence of spiking in the post-policy period is significantly lower, down to about 30% in the last two years of employment, with smaller declines in T-2 and T-3.

Next, we describe how the financial implications of salary spiking in TRSIL changed with the introduction of the 2005 excess compensation rule. These calculations focus on TRSIL models as providing an overall estimate of financial costs, including individual, institutional, and incidental sources of salary spiking , as described in Section 6.1. The models estimated and

[39] While CPTF is a natural group to consider as a counterfactual or falsification group, they appear to have substantially different pre-policy trends in both salary and spiking prevalence and confounding a potential causal interpretation of the effect of the policy.

reported in Table 3 for TRSIL identify approximately 2,395 employees as spiking per year, where at least one year from T to T-3 exceeds the 95% CI. The average FAS increase is $9,238 ($36,953 / 4). As noted by Fitzpatrick (2017), the state estimates that each dollar of FAS costs the pension system between $14 to $16, based on TRSIL assumptions about life expectancy and an 8.5% interest rate. Using the average value of $15 gives a total cost of $138,574 per individual with higher-than-expected compensation. Thus, in the pre-policy period, spiking may have cost $332 million per cohort of exiting employees. Starting in 2008, when the policy is fully implemented, only 1,717 employees are identified as spiking per year with an average FAS increase of $10,127 for a cost of $261 million. Thus, the policy may have reduced the impacts of salary spiking by approximately $71 million per year.[40]

Importantly, the figures discussed above do not include the revenue collected from excess compensation billings. This is somewhat complicated because only compensation increases above 6% are billed, and we do not have precise data on how much the state has collected from these billings. That said, there are reports that the state collected a cumulative $38 million by 2013-2014. Thus, computing a per-cohort cost might be about $3.8 million, a relatively modest amount compared to the suggested change in salary-spiking behavior. It should also be noted that this additional revenue represents a redistribution of resources to cover financial liabilities rather than a reduction in salary-spiking costs.

[40] Note that our calculation uses TRSIL estimates of future retirement benefits in a given school year and is not adjusted for inflation.

Figure 3. Compensation growth before and after excess compensation policy

Panel A: Pre-policy compensation growth

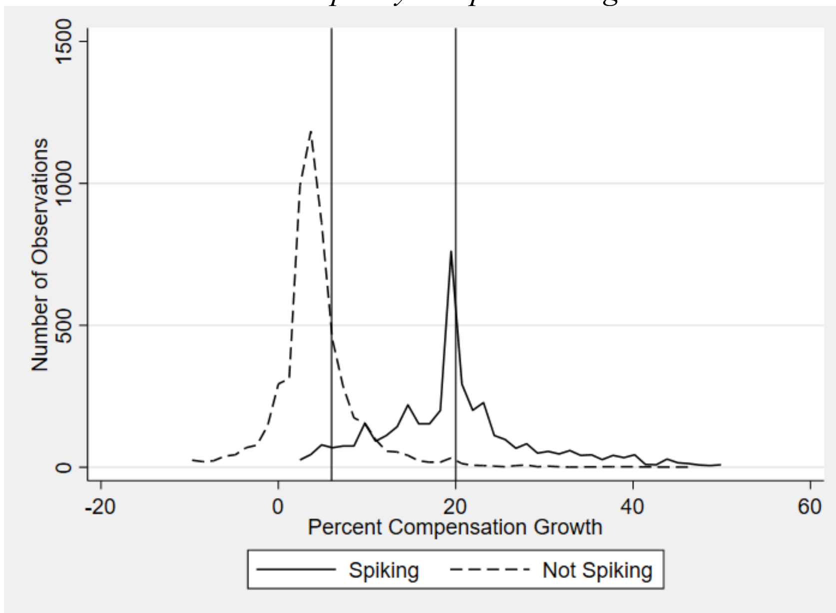

Panel B: Post-policy compensation growth

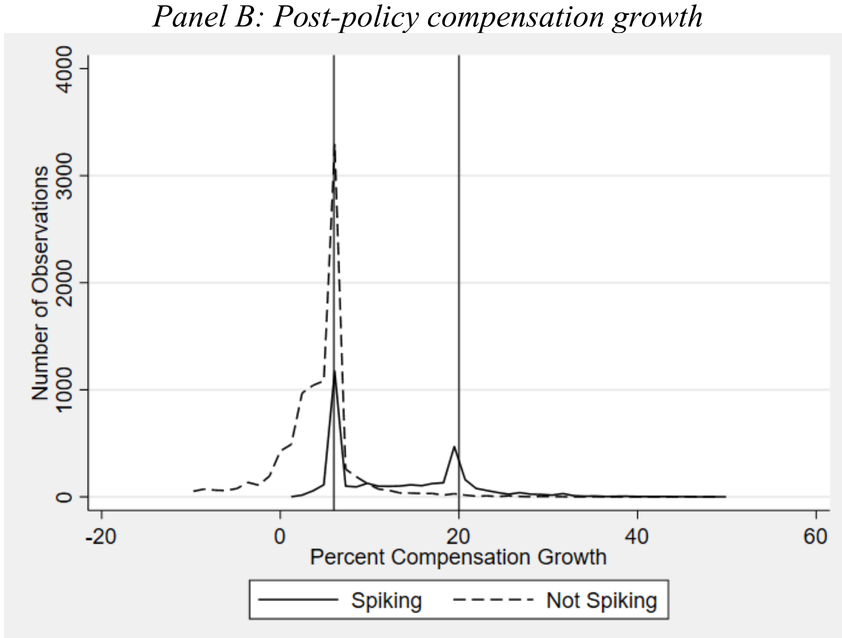

Notes: Each panel shows the distribution of compensation growth in a TRSIL employee's final year of employment by whether they are identified as spiking. Identification is based on the model from Table 2. The pre-policy period is prior to 2005, and the post-policy period is after 2007. Vertical lines indicate 6% and 20% compensation growth, where spiking policies are binding in post and pre-periods respectively.

We emphasize that this is only a rough calculation of financial costs. Assessing the precise financial implications in terms of pension wealth is challenging because we do not observe age or exact benefit factor for individuals, both of which are important factors for pension wealth calculations. For example, if individuals with higher than expected salaries tend to be older (and on average collect fewer benefits in retirement), the figures will overstate the true cost (and vice-versa). That said, one potential application of our approach could be use by state agencies with access to detailed information about age and benefit eligibility, allowing for more precise calculation of financial costs.

Lastly, in adopting the excess compensation rule, the state effectively established a definition of salary spiking: salary growth in the FAS period in excess of 6%. We contrast our method for identifying salary spiking with the state's definition of salary spiking in **Figure 3** by showing the distribution of compensation growth in an individual's final year of employment among employees we did, and did not, identify as spiking. Those who are identified as salary spiking are represented by the solid line and those who are not by dashed line. Panels A and B show results for the pre- and post-policy periods, respectively.

Two points stand out in **Figure 3**. First, this policy is targeted at employees who are salary spiking, and as such, there is a dramatic reduction in the number of employees near 20% growth in the post-policy period. Much of this mass is likely shifted to the new 6% level, both by reducing compensation growth among spikers, but also by transitioning from spiking to non-spiking. Second, the policy appears to have been unintended consequences for employees we do not identify as salary spiking. In the pre-policy period, a sizable portion of non-spiking employees have compensation growth above 6%, which is about a third of the distribution of non-spiking employees. This mass is greatly concentrated at 6% in the post-policy period.

Moreover, there is far less mass for non-salary spikers below 6%. This indicates that the decision to set excess compensation at 6% represents a balance of changing the behavior of employees who may and may not be salary spiking.

8. Conclusions

Salary spiking is a potential source of unfunded pension liabilities, but the prevalence and magnitude of spiking is not well understood. We develop an empirical definition of salary spiking that can be applied wherever reliable compensation data are available, and we use this method to identify the extent to which there appears to be salary spiking in Illinois.

Our analysis of finds that outside of Chicago Public Schools (which has its own pension system), about half of state pension-eligible employees are identified as salary spiking in their final year of employment. This contrasts with the findings from Chicago Public Schools, where spiking is only about 20% of pension-eligible employees are found to be salary spikers, but it is consistent with what we would expect given the incentives faced by school districts. School districts whose employees are members of the state pension plan (TRSIL) have little incentive to discourage salary spiking among their own employees since any unfunded liabilities generating by spiking are shared by all member districts. In contrast, Chicago Public Schools bares the full cost of any salary spiking that occurs in the district. This suggests that employer incentives play a key role in influencing salary spiking.

We also consider the implementation of an anti-spiking initiative instituted in Illinois in 2005. This initiative bills employers for liabilities associated with compensation growth in excess of 6%. Consistent with the notion that the excess compensation rule successfully internalized the costs of salary spiking across employers that had previously been shared across all member districts, we find large reductions in the prevalence of salary spiking after the rule went into effect. In particular, we estimate that spiking liabilities were reduced by $71 million per cohort of exiting employees. In fact, the prevalence of salary spiking in the post-policy period is somewhat similar to the levels of spiking identified in CTPF. However, we also find that policies that define spiking in terms of growth thresholds (such as the 2005 anti-spiking initiative) may have unintended consequences: employees with high salary growth who are not in fact salary spiking (e.g. consistently high salary growth each year) may have their compensation reduced due to the policy.[41]

This approach has important implications for policy. Understanding the ramifications of salary spiking is particularly relevant today given that states and localities are receiving a large, $122 billion, increase in funds connected with the American Rescue Plan.[42] If this spending goes to compensation that is pensionable, it could create unintended additional unfunded pension obligations. Moreover, this approach can be applied to any pension system where detailed compensation records are available as a method to explore individual cases of potential salary spiking.

That said, there is an important caveat to the findings presented above. Our approach identifies cases where salary deviates from prior patterns of compensation, but it is agnostic to the motivations behind the increase. As such, it is not possible to determine whether the salary spiking we observe is the result of employees seeking greater benefits from the pension system or simply incidental; for example, an employee may seek additional roles (such as mentorship) regardless of the implications for their pension. In fact, our analysis of non-exiting TRSIL employees indicates that a substantial portion of employees (about 20%) experience

[41] This need not be a cost, as some Illinois districts have been able to avoid excess compensation billings by using district CBAs to classify any increase in compensation above 6 percent as *non-pensionable*.
[42] https://www.ed.gov/news/press-releases/department-education-announces-american-rescue-plan-funds-all-50-states-puerto-rico-and-district-columbia-help-schools-reopen, accessed 4/19/2021.

unexpectedly large increases in compensation that are likely unrelated to DB pension benefits. This highlights the fact that employees frequently have discontinuous increases in compensation, not just near the end of their employment. Importantly, however, even if these increases are not *intended* to increase DB pension benefits, they affect pension benefits, and hence pension system liabilities, *regardless* of their intent.

References

Backes, B., Goldhaber, D., Grout, C., Koedel, C., Ni, S., Podgursky, M., ... & Xu, Z. (2016). Benefit or burden? On the intergenerational inequity of teacher pension plans. *Educational Researcher*, *45*(6), 367-377.

Biggs, Andrew G. (2015). *The State of Public Pension Funding: Are Government Employee Plans Back on Track?* Washington DC: American Enterprise Institute.

Brown, Jeffrey R., Robert Clark, and Joshua Rauh. (2011). "The Economics of State and Local Pensions." *Journal of Pension Economics and Finance* 10(2):161–72.

Costrell, Robert M; Maloney, Larry. (2013). *The Big Squeeze: Retirement Costs and School District Budgets*. Washington D.C.: Thomas B Fordham Institute.

Costrell, Robert M., Collin Hitt, and James V Shuls. (2020). "A $19 Billion Blind Spot: State Pension Spending." *Educational Researcher* Published online April 15, 2020.

Costrell, Robert M., and Michael Podgursky. (2009). "Peaks, Cliffs, and Valleys: The Peculiar Incentives in Teacher Retirement Systems and Their Consequences for School Staffing." *Education Finance and Policy* 4(2):175–211.

Fitzpatrick, Maria D. (2015). "How Much Are Public School Teachers Willing to Pay for Their Retirement Benefits." *American Economic Journal: Economic Policy* 7(4):165–88.

Fitzpatrick, Maria D. (2017). "Pension-Spiking, Free-Riding, and the Effects of Pension Reform on Teachers' Earnings." *Journal of Public Economics* 148:57–74.

De Gooijer, Jan G., and Rob J. Hyndman. (2006). "25 Years of Time Series Forecasting." *International Journal of Forecasting* 22(3):443–73.

Goldhaber, D., Grout, C., and Holden, K., (2018). *Public Pensions and Salary Spiking: A Cautionary Tale of Data Inaccuracy Leading to Erroneous Results.* CALDER Working

Paper No. 202-0918-1.

Greene, William H. (2003). *Econometric Analysis*. 5th ed. edited by R. Bannister, P. J. Boardman, G. Soto, and M. McHale. Upper Saddle River, NJ: Pearson Education, Inc.

Hyndman, Rob J., Anne B. Koehler, Ralph D. Snyder, and Simone Grose. (2002). "A State Space Framework for Automatic Forecasting Using Exponential Smoothing Methods." *International Journal of Forecasting* 18:439–54.

Makridakis, Spyros, and Michèle Hibon. (1987). "Confidence Intervals: An Empirical Investigation of the Series in the M-Competition." *International Journal of Forecasting* 3:489–508.

Malanga, Steven, and Josh McGee. (2018). *Garden State Crowd-Out: How New Jersey's Pension Crisis Threatens the State Budget*. New York , NY.

Mannino, Michael V, and Elizabeth S. Cooperman. (2013). "Benefit Enhancement in Public Employee Defined Benefit Pension Plans: Evidence from Three Sources." University of Colorado Denver Working Paper.

Mincer, Jacob. (1974). *Schooling, Experience, and Earnings*. New York: Columbia University Press.

Moe, Terry M. (2006). "The Union Label on the Ballot Box." *Education Next* 6(3):58–67.

Murphy, Kevin M., and Finis Welch. (1990). "Empirical Age-Earnings Profiles." *Journal of Labor Economics* 8(2):202–29.

Newbold, P., and C. W. J. Granger. (1974). "Experience with Forecasting Univariate Time Series and the Combination of Forecasts." *Journal of the Royal Statistical Society* 137(2):131–65.

Novy-Marx, Robert, and Joshua Rauh. (2011). "Public Pension Promises: How Big Are They and What Are They Worth?” *Journal of Finance* 66(4):1211–49.

Novy-Marx, Robert, and Joshua D. Rauh. (2009). “The Liabilities and Risks of State-Sponsored Pension Plans.” *Journal of Economic Perspectives* 23(4):191–210.

Shuls, James V, and Joseph Lux. (2019). “Examining the Prevalence of Salary Spiking in Missouri’s Defined Benefit Public School Retirement System James V. Shuls 1 and Joseph Lux 2 1.” *International Review of Accounting, Banking and Finance* 11(1):1–11.

Wallis, Kenneth F. (1974). “Discussion of ‘Experience with Forecasting Univariate Time-Series and the Combination of Forecasts’” by P. Newbold and CWJ Granger. JR.” *Journal of the American Statistical Association* Series A(137):155–56.

Williams, W. H., and M. L. Goodman. (1971). “A Simple Method for the Construction Forecasts for Economic Limits Confldence.” *Journal of the American Statistical Association* 66(336):752–54.

Appendix A. Additional tables and figures

Table A1. Prevalence and magnitude of salary spiking with 80% CI for TRSIL

	(1)	(2)	(3)	(4)
Model Fit (Avg. R-squared)	0.92	0.93	0.93	0.93
Coefficient on quadratic term (Avg.)	66.11	31.83	15.38	-0.19
Proportion Spiking in				
T (Final year)	0.50	0.57	0.58	0.58
T-1		0.47	0.52	0.53
T-2			0.36	0.42
T-3				0.33
Magnitude among spiking (Avg.)	$11,351	$21,851	$30,185	$41,594
Observations	46,561	46,561	46,561	46,561

Notes: This table results for TRSIL employees who exit employment using an 80% CI interval instead of 95% as shown in Table 2. Model fit is summarized by adjusted R-squared terms. Each regression includes school year and school year squared, and average coefficients on the quadratic term are reported in the second row. Prevalence is a sum over an indicator for whether an individual has a salary exceeding the CI given by equation (3). The total magnitude of salary spiking is measured as the sum over difference $C_{iT} - \hat{C}_{iT}$ *for each year an individual is identified as spiking. Samples come from a panel of employment data from 1991-1992 to 2011-2012, limited to employees with at least 10 years of observable consecutive employment at the end of their careers.*

Table A2. Prevalence and magnitude of salary spiking

	(1)	(2)	(3)	(4)
Panel A: TRSIL				
Model Fit (Avg. R-squared)	0.92	0.93	0.93	0.93
Coefficient on quadratic term (Avg.)	66.11	31.83	15.38	-0.19
Proportion Spiking in				
T (Final year)	0.36	0.46	0.48	0.46
T-1		0.33	0.39	0.40
T-2			0.19	0.26
T-3				0.15
Magnitude among spikers (Avg.)	$13,431	$24,221	$32,263	$43,500
Observations	46,561	46,561	46,561	46,561
Panel B: CTPF				
Model Fit (Avg. R-squared)	0.90	0.92	0.92	0.91
Coefficient on quadratic term (Avg.)	72.88	80.29	85.00	69.98
Proportion Spiking in				
T (Final year)	0.15	0.19	0.20	0.20
T-1		0.21	0.24	0.26
T-2			0.16	0.20
T-3				0.17
Magnitude among spiking (Avg.)	$11,583	$17,548	$24,622	$37,519
Observations	13,163	13,163	13,163	13,163
Panel C: Non-exiting TRSIL employees				
Model Fit (Avg. R-squared)	0.91	0.91	0.91	0.90
Coefficient on quadratic term (Avg.)	42.73	46.36	39.57	32.66
Proportion Spiking in				
T (Final year)	0.10	0.15	0.19	0.21
T-1		0.11	0.17	0.20
T-2			0.12	0.18
T-3				0.13
Magnitude among spiking (Avg.)	$9,019	$16,935	$30,988	$48,701
Observations	26,880	26,880	26,880	26,880

Notes: Panel A reports results for TRSIL employees who exit employment, and Panel B reports results for exiting CTPF employees. Panel C presents results using non-exiting TRSIL individuals who are observed in all years of the data, and a randomly selected *pseudo-final year. Model fit is summarized by adjusted R-squared terms. Each regression includes school year and school year squared, and average coefficients on the quadratic term are reported in the second row. Prevalence is a sum over an indicator for whether an individual has compensation exceeding the CI given by equation (3). The total magnitude of salary spiking is measured as the sum over difference* $C_{iT} - \hat{C}_{iT}$ *for each year an individual is identified as spiking.*

Figure A.1 Residuals according to year until exit, by fitted and forecast

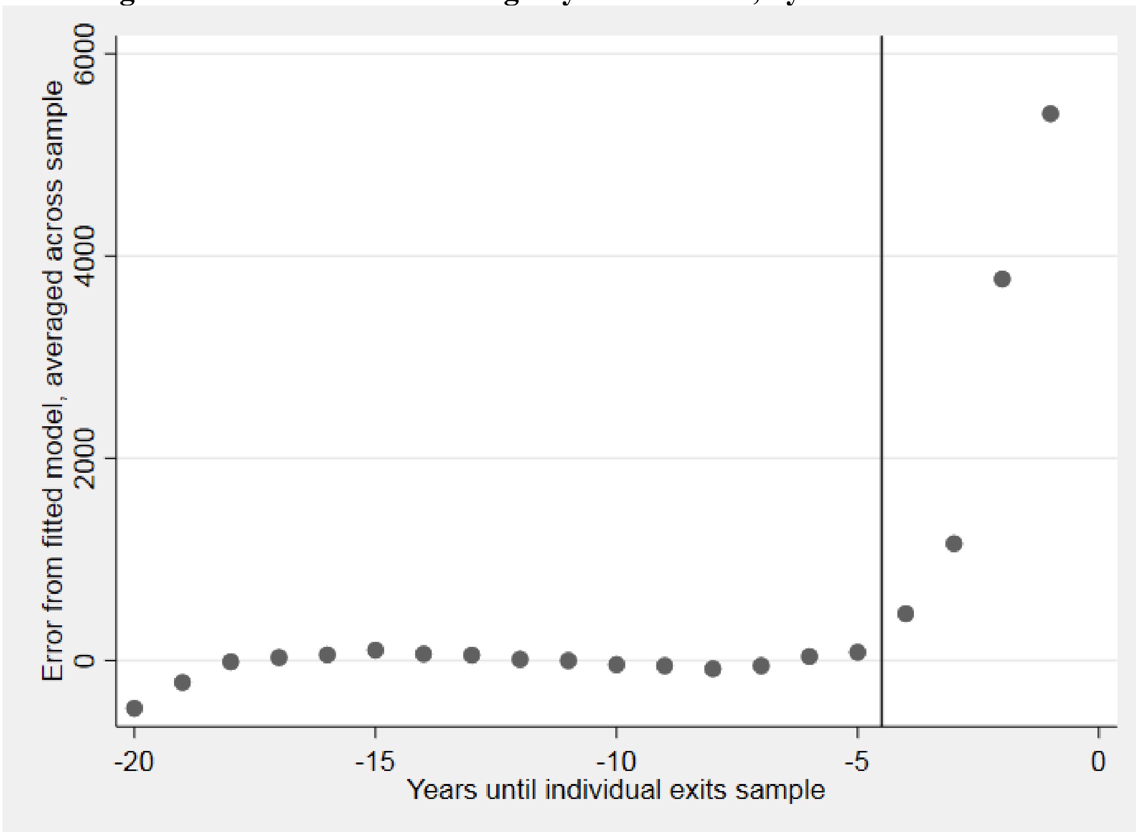

Notes: The figure indicates residuals from the model presented in Table 1, column (4). Each point represents the average residual for a given year of employment relative to their separation across all individuals in the sample. The vertical line indicates that points to the left are fit to the model, while points to the right compare salary to forecasted salary.