



Istituto Nazionale Previdenza Sociale

aprile 2026 – numero 119



WorkINPS *Papers*

## **Salary caps in the Public Administration**

**Daniele Checchi  
Francesco Figari  
Carlo Fiorio**

**ISSN 2532 -8565**

Lo scopo della serie WorkINPS papers è quello di promuovere la circolazione di documenti di lavoro prodotti da INPS o presentati da esperti indipendenti nel corso di seminari INPS, con l'obiettivo di stimolare commenti e suggerimenti.

Le opinioni espresse negli articoli sono quelle degli autori e non coinvolgono la responsabilità di INPS.

The purpose of the WorkINPS papers series is to promote the circulation of working papers prepared within INPS or presented in INPS seminars by outside experts with the aim of stimulating comments and suggestions.

The views expressed in the articles are those of the authors and do not involve the responsibility of INPS.

Responsabile Scientifico

**Agar Brugiavini**

Comitato Scientifico

**Agar Brugiavini, Vito La Monica, Gianfranco Santoro.**

*In copertina: uno storico "Punto cliente" a Tuscania  
INPS, Direzione generale, Archivio storico*

## I WORKINPS PAPER

Le basi dati amministrative dell'*INPS* rappresentano una fonte statistica unica per studiare scientificamente temi cruciali per l'economia italiana, la società e la politica economica: non solo il mercato del lavoro e i sistemi di protezione sociale, ma anche i nodi strutturali che impediscono all'Italia di crescere in modo adeguato. All'interno dell'Istituto, questi temi vengono studiati sia dai funzionari impiegati in attività di ricerca, sia dai *VisitInps Scholars*, ricercatori italiani e stranieri selezionati in base al loro curriculum vitae e al progetto di ricerca presentato.

I **WORKINPS** hanno lo scopo di diffondere i risultati delle ricerche svolte all'interno dell'Istituto a un più ampio numero possibile di ricercatori, studenti e policy markers.

Questi saggi di ricerca rappresentano un prodotto di avanzamento intermedio rispetto alla pubblicazione scientifica finale, un processo che nelle scienze sociali può chiedere anche diversi anni. Il processo di pubblicazione scientifica finale sarà gestito dai singoli autori.

**Agar Brugiavini**

# Salary caps in the Public Administration

**Daniele Checchi**

(Università degli Studi di Milano)

**Francesco Figari**

(Università degli Studi del Piemonte  
Orientale)

**Carlo Fiorio**

(Università degli Studi di Milano  
and Irvapp-FBK)

## **Tetti salariali nella Pubblica Amministrazione**

Daniele Checchi, Francesco Figari, Carlo Fiorio

### **Abstract**

Studiamo gli effetti dei tetti salariali per i dirigenti pubblici apicali in Italia, con particolare riferimento alle riforme del 2011 e del 2014 introdotte durante il periodo di consolidamento fiscale successivo alla crisi. Il primo tetto, fissato a 290.000 euro nel 2011 e successivamente ridotto a 240.000 euro nel 2014, ha imposto un'aliquota marginale del 100% sui redditi superiori alla soglia. Utilizzando dati amministrativi collegati impresa-lavoratore e un disegno di event study difference-in-differences, stimiamo l'impatto causale su retribuzioni e partecipazione.

La riforma del 2011 ha significativamente compresso i salari nella parte alta della distribuzione del settore pubblico e ha generato consistenti risparmi fiscali.

Le risposte sul margine intensivo sono modeste: l'elasticità dei redditi imponibili netti nella fascia più alta è pari a circa 0,214, e la conseguente perdita di efficienza per chi è rimasto in servizio è contenuta, riflettendo distorsioni di secondo ordine nel triangolo di Harberger.

Al contrario, le risposte sul margine estensivo sono rilevanti: la riforma ha aumentato significativamente il pensionamento anticipato e il passaggio al settore privato, con elasticità della partecipazione di gran lunga superiori all'elasticità sul margine intensivo. L'analisi di benessere mostra che i costi complessivi di efficienza sono determinati principalmente da questi margini estensivi.

Parole chiave: tetti salariali; pubblica amministrazione; Italia

Keywords: salary cap; public administration; Italy

# Salary caps in the Public Administration

Daniele Checchi\*, Francesco Figari†, Carlo Fiorio‡

## Abstract

We study the effects of salary caps for top public managers in Italy, focusing on the 2011 and 2014 reforms implemented during the post-crisis fiscal consolidation period. The first cap, set at €290,000 and later reduced to €240,000, imposed a 100% marginal tax rate above the threshold. Using linked employer–employee administrative data and an event-study difference-in-differences design, we estimate the causal impact on earnings and participation. The 2011 reform significantly compressed salaries at the top of the public-sector distribution and generated sizable fiscal savings.

Intensive-margin responses are modest: the top-bracket elasticity of taxable net earnings is about 0.214, and the resulting efficiency loss for stayers is small, reflecting second-order Harberger-triangle distortions. In contrast, extensive responses are large: the reform significantly increased early retirement and switching to the private sector, with participation elasticities far exceeding the intensive elasticity. Welfare analysis shows that overall efficiency costs are driven primarily by these extensive margins.

---

We are grateful to Giulio Casilli for excellent assistance in interpreting Italian law, Stefania Albanesi, Erich Battistin, Francesco Fasani, Maxime Gravouelle, Wouter Leenders, Enrico Rubolino, and Ben Sprung-Keyser for useful discussions; and participants at the XXXVII Italian Public Economics Society (SIEP) Conference (Naples), the 118th National Tax Association Annual Conference (Boston), the 66th Italian Economic Society (SIE) Conference (Naples), and the VisitINPS online seminar for helpful comments and suggestions. The first version of this project was presented online in a Visitinps seminar on July 18th, 2023. The data used in this paper are property of the Italian Social Security Institute (INPS) and are accessible by researchers at the INPS premises through the VisitInps program (see <https://www.inps.it/it/it/dati-e-bilanci/attivita-di-ricerca/programma-visitinps-scholars.html>). To access data for replication purposes, researchers should contact the INPS Research department ([dcstudiericerche@inps.it](mailto:dcstudiericerche@inps.it)). The findings and conclusions expressed are solely those of the author and do not represent the views of INPS.

\*Università degli Studi di Milano. [daniele.checchi@unimi.it](mailto:daniele.checchi@unimi.it)

†Università degli Studi del Piemonte Orientale: [francesco.figari@uniupo.it](mailto:francesco.figari@uniupo.it)

‡Università degli Studi di Milano and Irvapp-FBK. [carlo.fiorio@unimi.it](mailto:carlo.fiorio@unimi.it)

# 1 Introduction

The aftermath of the 2008 global financial crisis triggered profound changes in public-sector pay setting across Europe. Governments faced acute pressures to contain public expenditures while maintaining state capacity and the ability to attract and retain skilled managers. Cross-country studies from this period document significant public-private pay differentials, with average public-sector compensation often exceeding private-sector levels, though much of this gap is explained by differences in worker characteristics such as age, education, and occupation (Giordano et al., 2011; Campos et al., 2017). In response to mounting fiscal imbalances, many European countries implemented austerity packages combining salary freezes, reductions in personnel, and limits on salary progression (Caballero-Cordero et al., 2025). These measures compressed public pay structures and reshaped managerial incentives, prompting concerns over equity, efficiency, and the long-term attractiveness of public service (OECD, 2012). A salient component of these reforms was the introduction of salary controls—fixed or benchmark-relative caps on earnings at the top of the public-sector distribution.

Across Europe, salary-cap reforms shared common objectives: imposing discipline on the public wage bill, reducing disparities between top and bottom earners, and increasing transparency. Although the precise calibration of caps and their scope varied, the central rationale was to restore fiscal sustainability while signaling fairness amid widespread cuts (Bruni, 2017). Yet the behavioral and welfare consequences of such policies remain poorly understood. By directly constraining the upper tail of the public pay distribution, salary caps create sharp changes in effective marginal tax rates and thereby offer a rare opportunity to study labor-supply responses at the top of the income distribution in an institutional environment where avoidance or bargaining channels are limited.

This paper studies the economic and welfare consequences of salary caps for top public managers in Italy. Following the sovereign debt crisis, the Italian government introduced a statutory ceiling on top public-sector earnings: in 2011, the cap was set at the salary of a Supreme Court judge (€290,000) and was later reduced to €240,000

in 2014. Earnings above the cap face a 100% marginal tax rate, implying that any salary above the threshold is fully confiscated. The government justified these policies as part of a cost-containment strategy, and Constitutional Court rulings reinforced their legitimacy, affirming that they are valid tools for ensuring predictable and stable public-resource planning. However, the cap does not apply to constitutional bodies or to top managers of state-owned enterprises organized as joint-stock companies, where competitive pressures from the private sector are deemed stronger. These institutional features generate rich variation in exposure to the policy, which we exploit in our empirical analysis.

We study these reforms using administrative data from the Italian Social Security Institute (INPS), which provides anonymized employer-employee-matched records covering the entire universe of salaried employment and allows us to track public managers over time. These data contain precise information on earnings, employment spells, and identifiers for both employers and employees, making it possible to merge with additional demographic and job-related characteristics. Because exemptions to the cap apply only to specific institutions, we compare managers in capped organizations to managers in institutions not subject to the cap but with identical hiring rules and career structures. This design allows us to construct a credible control group and estimate the policy’s intention-to-treat effects.

Our analysis connects to the modern public-finance literature on top incomes, which studies how high-salary individuals respond to large changes in effective marginal tax rates. Behavioral responses at the top may involve adjustments in labor supply, avoidance, bargaining, or participation (Piketty et al., 2014). Salary caps generate a particularly stark setting for studying these mechanisms: any salary above the cap is taxed at a 100% marginal rate, eliminating the return to effort on inframarginal tasks. We adopt the sufficient-statistics framework for welfare analysis, which links deadweight losses to observed elasticities around kinks or discrete jumps in net-of-tax rates (Saez, 2010; Chetty, 2009). Although salary is truncated by the cap, the methodology remains valid because behavioral responses can be expressed in terms of the pre-cap marginal net-of-tax rate. At the top of the Italian public-sector pay distribution, avoidance channels are limited by statute, and hours are essentially

non-contractible, making the distinction between intensive and extensive margins especially important. Against this backdrop, transitions to the private sector or early retirement can become the dominant behavioral channel.

Empirically, we exploit the staggered implementation of cap policies to estimate their impact on earnings, participation decisions, and welfare. Our results are three-fold. First, salary caps achieved their primary fiscal objective. Affected individuals experience an average reduction in gross earnings of over €130,000 per year. When scaled to the full population, aggregate gross savings exceed €74 million per year, while net-of-tax savings—recognizing that forgone earnings would otherwise have been taxed at 43%—amount to about €42 million. These figures underscore the fiscal significance of the reform and the importance of accounting for both gross and net effects.

Second, intensive-margin responses are modest. Despite a large drop in the marginal net-of-tax rate from 0.57 to zero, we estimate a taxable-income intensive elasticity of  $\approx 0.214$ . Using standard calibrations of effort disutility, we compute an intensive-margin welfare loss of roughly 12,250 per affected individual, which is the reduction in private surplus net of saved effort costs, and is much lower than the full mechanical income loss.

Third, extensive-margin responses dominate both behavioral and welfare adjustments. The probabilities of switching to private-sector employment and of early retirement rise sharply after the 2011 reform. The corresponding transition elasticities— $\varepsilon^{\text{switch}} \approx 1.41$  and  $\varepsilon^{\text{retire}} \approx 5.29$ —are an order of magnitude larger than the intensive-margin elasticity, indicating that the primary behavioral response occurs through exit rather than reductions in labor supply among stayers. Consistent with the public-finance literature on large fiscal discontinuities, these extensive-margin elasticities exceed one because the baseline probabilities of switching or retirement are low, and individuals face a discrete set of occupational alternatives (e.g. [Kleven et al., 2013](#); [Kleven and Waseem, 2013](#); [Piketty et al., 2014](#)). Under a planner-centric benchmark, extensive-margin responses generate an additional welfare loss of roughly €12,700 per eligible manager, raising total per-person losses to about €24,950. These effects highlight the centrality of participation margins in assessing the overall effi-

ciency cost of salary caps. Although we cannot observe managerial output directly, our results suggest that the reallocation of high-ability public managers away from the public sector may carry broader implications for state capacity. Exit decisions may be delayed as managers wait for suitable outside options, meaning that the full effects of a cap may not materialize immediately. As recent research has emphasized (Kleven et al., 2025), such delayed yet discrete adjustments are a key channel through which high-income individuals respond to strong pay compression. Because top public managers influence organizational performance, any induced reallocation or early exit may have consequences beyond the wage bill itself.

Our study complements and extends Di Porto et al. (2025), who analyze the 2014 Italian salary cap. In contrast, we show that the 2011 reform produced the largest and most persistent effects on both earnings and participation margins. Moreover, our analysis uses a representative INPS dataset covering the entire employee population, in contrast to their non-representative sample of public-sector contributors. We further argue that their chosen control group is affected by substantial compositional changes over time, due to the fixed-term nature of public managerial assignments. By instead comparing capped managers to managers in exempt institutions with identical recruitment rules, we obtain more credible intention-to-treat estimates. While both studies find that caps increase exit from public employment, our results show that early retirement is an important additional margin of response.

The remainder of the paper is organized as follows. Section 2 presents the institutional setting. Section 3 describes the data, and Section 4 outlines our empirical strategy. Section 5 presents the main results, and Section 6 discusses their implications for elasticities and social welfare. Section 7 concludes.

## 2 Institutional Setting

Italy introduced salary caps for top public managers in the aftermath of the 2008 financial crisis and the 2011 sovereign debt crisis. In 2011, the government set a ceiling on gross annual remuneration at the salary of a Supreme Court judge (€290,000). Earnings above this threshold were subject to a 100% marginal tax rate, implying full confiscation of salary beyond the cap. In 2014, the ceiling was lowered to €240,000

as part of a broader reform titled “For a Moderate Public Administration.” The stated objective of both measures was to contain costs by reducing and rationalising the wage bill at the upper end of the public pay distribution.

The salary cap applies to top managers in the central public administration, including ministries, agencies, and regulatory authorities. However, important exemptions remain. Constitutional bodies—including the Presidency of the Republic, Parliament, the Senate, and the Constitutional Court—retain autonomy over remuneration and are not directly bound by the cap. Managers in state-owned enterprises (SOEs) structured as joint-stock companies are also exempt, reflecting the view that such positions compete with the private sector for talent and require greater pay flexibility.

Since the 1990s, a distinctive feature of public-sector careers in Italy has been the widespread use of temporary task-based managerial assignments. These are not permanent managerial ranks but positions linked to specific responsibilities and associated salary supplements, typically lasting three to five years. Although formally temporary, in practice many executives continue to hold highly paid managerial roles over time by renewing an existing assignment, moving to comparable roles in other public administrations, or, in some cases, accumulating multiple assignments (and the corresponding additional payments), including across different public-administration bodies. It remains possible, though less common, for managers to return to lower-level positions, which would entail a substantial reduction in pay. As a result, cross-sectional comparisons of workers above and below the cap may be biased by endogenous changes in composition. Our empirical strategy, therefore, leverages exempt institutions — those that apply identical hiring, promotion, and assignment rules but are not subject to the cap — to construct a credible control group.

Judicial review has played a role in shaping the evolution of salary caps. The Constitutional Court has upheld the legitimacy of caps as tools to ensure the stability of public finances and promote equity in remuneration. Nonetheless, recent rulings have questioned the rigidity of the 2014 fixed cap, suggesting the need for dynamic,

benchmark-linked ceilings that can evolve with economic conditions.<sup>1</sup> Despite these developments, the core structure of the cap remains in place, and exemptions continue to generate useful variation for empirical analysis.

Overall, the Italian institutional setting provides a compelling environment for studying behavioural responses to large reductions in net-of-tax rates at the top of the public pay distribution. The combination of statutory rigidity, limited avoidance channels, heterogeneous exposure across institutions, and fixed-term managerial assignments creates a natural setting to isolate the intensive and extensive margins of adjustment that motivate our empirical and welfare analysis.

### 3 Data

Our data are provided by the Italian Social Security Institute (*Istituto Nazionale della Previdenza Sociale*, INPS), which manages contributions, unemployment benefits, layoffs, and pension payments. INPS makes available its anonymised administrative data through a remote access system (VisitInps program). The administrative archives provided within the VisItinps program, are the only data currently available to researchers to assess the effects of the salary caps introduced in 2011 and 2014 in Italy, as they come as longitudinal data, hence following individual workers as they enter the INPS archives, with detailed information regarding each social contribution payment, the related salary, duration of the employment spell, identifier of the employer and so on. The INPS administrative archives comprise employer–employee-matched records that link each social security contribution to the firm responsible for payment and the worker on whose behalf the contribution was made.

On INPS servers, there are three files potentially useful for studying the impact of the 2011 and 2104 reforms: (a) *Privati Uniemens* (Uniemens hereafter), which contains the population of employees who contribute to the private sector pension fund (*FPLD-Fondo Pensioni Lavoratori Dipendenti*), excluding employees in agriculture and caregivers. The dataset also includes a limited group of public-sector employees (e.g. INPS employees), whose inclusion reflects historical institutional ar-

---

<sup>1</sup>See Supreme Court sentence no. 135/2025.

rangements. These workers originated from company-based pension funds that were gradually merged into the general private-sector pension fund, primarily consist of employees of state-owned enterprises, and are mostly exempt from the earnings cap. (b) *Pubblici Uniemens* (INPDAP hereafter), containing the population of employees that contribute to the public pension fund (previously known as *INPDAP-Istituto Nazionale di Previdenza e assistenza per i Dipendenti dell'Amministrazione Pubblica*). This archive is made available since 2014; (c) *Estratti Conto* (EC hereafter), which focuses on workers born in 1950 or later who contributed to any fund managed by INPS. Since even self-employed individuals are obliged to contribute to the national pension system, it covers almost the entire workforce of the country, except for workers in the informal sector and professionals registered with professional bodies (such as lawyers and accountants). This is constructed as a 13.2% random sample of all workers, selected by date of birth, using four fixed days in each month over the entire period considered.

All INPS archives share some major limitations. First, they have no information on workers' educational attainment; however, this information is mostly time-invariant after labour market entry. In addition, we focus on the very top of the salary distribution, where individuals without a tertiary degree are unlikely to be found. Second, because INPS's primary aim is to collect social security contributions from employers and employees and to manage payments for pensions and other benefits, these payments are made on an accrual basis rather than on a cash basis, which provides a noisy measure of annual salaries.<sup>2</sup>

The INPDAP archive would be the ideal source of information were it not for the fact that no records are available before 2014, which is why we exclude them from the current analysis. The Uniemens archive includes only a subset of the contributions paid by public-sector employers and therefore is not representative of the

---

<sup>2</sup>The salary measure is therefore contaminated by retroactive payments, which can be large (especially in the occurrence of contractual renewals and/or when performance-related components are paid - typically in the spring of the ensuing year): these payments are not relevant for the policy implementation, as the salary cap is introduced with a definition of salary on a cash basis over the calendar year. This implies that INPS data indicate that some salaries were paid above the cap, even among public administration entities subject to it. Over the years, these differences have been compensated for, but noise in annual salaries remains.

public-sector universe, on which we focus.

As a consequence, we use the EC archive, which randomly selects all workers irrespective of the pension fund they contributed to, thus consisting of a dataset that is representative of the whole population of workers born in 1950 or later, and discuss further the issue of the choice of the INPS data in Appendix A. By leveraging anonymised unique identifiers for both employers and employees, these data can be merged with additional information on their individual characteristics. We disentangle employers belonging to public administrations (PAs) from private-sector employers, using an archive of the universe of PAs publicly provided by the Ministry of Economy and Finance.<sup>3</sup> This dataset includes the name of each entity, its public-sector category, additional administrative information that distinguishes between public administrations subject to the salary cap and those exempt from it, and its unique tax code identifier. INPS encrypted this tax code identifier and matched it to its archives, allowing us to augment the information available in the EC data (the employer’s public status, its main activity, and whether it is exempt from the salary cap).

Given the nature of the EC dataset, in 2019, we have a representative sample of workers aged 69 or younger, representing nearly the entire universe of public sector employees.<sup>4</sup> After adjusting for the random sampling rate (48 birth dates per year), the EC data cover approximately 3.9 million public workers. Approximately 300,000 of these are employed by organisations exempt from the salary cap. The dataset includes nearly 13,000 institutions, of which 56 are exempt. These aggregates closely match official statistics, whereas reliance on Uniemens data would result in a highly biased representation of public-sector employment (Table 1).

---

<sup>3</sup>The list of all PAs is available here (<https://bdap-opendata.rgs.mef.gov.it/content/anagrafica-enti-ente>). We downloaded the file on 31 March 2023.

<sup>4</sup>In Italy, the ordinary statutory retirement age for public employees is 67 years. In limited, specific cases, the law allows a temporary continuation of service up to age 70, mainly to meet organisational or institutional needs. This possibility has historically applied to certain sectors, such as academia (e.g., university professors), and, under previous regulations, to the judiciary, although these extensions have been progressively restricted by recent reforms. In any case, employment as a public-sector employee cannot continue beyond the age of 70, after which only non-employee roles (e.g., consultancy or advisory positions) may be permitted.

EC records all types of work contracts and matches employees to employers. Therefore, we can construct information on each worker’s employment status and salary for each year and workplace. Because workers can have multiple contracts in a year, we identify the relevant employer for the year as the one that pays the highest salary. Because public-sector workers can occasionally work for private firms upon obtaining authorisation from their primary employer, we focus only on payments received from a PA but also track all wages and salaries from private firms, building a panel dataset of individual yearly wages and salaries.

In Table 2, we present some descriptive statistics contrasting the PA with the private sector (columns A and B), including all records in the EC sample. We observe that in Italy, the PA share of women is 60% compared to 41% in the private sector, the average age is 8 years older, and the share of foreign workers is marginal, compared to 15% in the private sector. This reflects the recruitment rules in the PA, which require an open competition, with language proficiency representing a significant hurdle for foreign-born citizens.

Mean salaries in the PA are larger than in the private sector, which reflects the fact that a PA pays its employees mostly according to salaries established by collective labour contracts (10th percentile salary is five times larger than the corresponding salary in the private sector), that PA is less likely to hire workers in short-term or few hours-per-week contracts, more likely offering permanent contracts, and that access to the public sector is gained by a public competition that requires minimum level of education (large part of PA employees are required to hold a tertiary degree, e.g. in the healthcare, university or ministries, including the education sector).

Table 2 allows us to explain three further major choices we adopted for constructing our final sample. First, given the nature of the reforms we consider, which affect only top earners, we excluded workers who are very unlikely to be affected by the policy during their careers. To this aim, as the 99.5 quantile of earnings in the whole sample in 2019 is equal to € 123,611 and corresponds to about € 100,000 in 2005, using the consumer price index (CPI), we identify all workers that over the period 2005-2019 earned at least once a yearly salary above € 100,000 (in 2005 prices, updated over the years using the CPI). Second, because we focus on the

2005-2019 period, just before the COVID-19 pandemic, and the EC dataset is limited to individuals born in 1950 or later, we restricted the analysis to workers born in 1950-1970. This allows us to follow the cohort of workers who, at the start of our period, were aged 35-55 and to track them as they aged 49-69 in 2019, just before the outbreak of the COVID-19 pandemic. Finally, because salary cap policies apply only to workers in the public administration and the broader public sector, and because public-sector workers differ systematically from private-sector workers in motivation, risk attitudes, and other characteristics,<sup>5</sup> as in [Di Porto et al. \(2025\)](#), we restrict the sample to public-sector workers.

## 4 Empirical strategy

To identify the causal effect of the salary cap policies introduced in 2011 and in 2014 on PA workers, we define two mutually exclusive groups of eligible workers. Having selected all workers in our sample who ever worked in the PA and earned € 100,000 (at 2005 prices) or more for at least one year in the period 2005-2019, we divided the sample into workers employed by PAs subject to or exempt from salary caps. Treated individuals belong only to the first group. However, even among salary-capped PAs, eligible individuals are not univocally identified for two main reasons. First, the salary variable recorded in all INPS archives measures earnings for the calculation of future pensions, including arrears and accrual-basis definitions, not on a cash basis, as used for the implementation of the salary cap. This means that none of the INPS archives show a clear salary cut of € 290,000 and € 240,000 for the 2011 and 2014 policies, respectively. Second, because policies are implemented over several years, they are likely to affect individuals who did not reach the cap at the time of implementation but would have if no salary cap had been introduced.

Hence, we define two mutually exclusive eligibility groups based on pre-policy salary thresholds. The first group comprises individuals who ever had an annual

---

<sup>5</sup>A large literature documents that public sector employees systematically differ from private sector workers in underlying motivations and preferences, including higher levels of public service motivation, pro-social attitude and different responses to incentives, as well as distinct risk attitudes and job-selection behavior (e.g., see [Perry and Wise, 1990](#); [Rainey and Steinbauer, 1999](#); [Buelens and den Broeck, 2007](#); [Bellante and Link, 1981](#); [Buurman et al., 2012](#); [Fenizia et al., 2023](#)).

earning above € 225,000 during the period 2005–2010. The second group includes individuals who ever had an annual salary above € 175,000 during the period 2005–2013, provided they were not included in the first eligibility group. These eligibility rules are based on earnings before each policy’s implementation and remain fixed throughout the analysis. This definition reflects that, based on pre-policy earnings alone, we cannot perfectly identify who would have been affected by the salary cap after its implementation.

As shown in Table 3, individuals closer to the cap are more likely to surpass it within five years. Moreover, reaching higher ranks is not permanent: only 56.5% of those earning more than € 290,000 remain in the same earning group five years later, reflecting the fixed-term nature of managerial assignments.

Restricting eligibility to those who crossed the cap before the policy would incorrectly classify many individuals who would have crossed it later as controls.<sup>6</sup> Our approach, therefore, estimates the effect of eligibility rather than the effect of actual treatment.

The final sample of observations we use to estimate the effects of salary caps for top-earning employees in the PA is an unbalanced panel of 93,673 observations over 14 years, of which 3.6% are defined as eligible for the 2011 policy and 8.1% for the 2014 policy (Table 4). Controls are all workers employed in the PA, exempt from the salary cap, and workers who work in salary-capped PAs but are not considered eligible under the eligibility rule defined above. Some robustness checks on the eligibility rule are investigated in Section 5.2.

We then estimate the following event-study difference-in-differences model:

$$\begin{aligned}
 y_{it} = & \alpha_i + \delta_t + \sum_{\tau \neq 2010} \beta_{\tau}^{2011} (\mathbb{1}\{\text{Year}_t = \tau\} \times \text{Eligible}_i^{2011}) \\
 & + \sum_{\tau \neq 2010} \beta_{\tau}^{2014} (\mathbb{1}\{\text{Year}_t = \tau\} \times \text{Eligible}_i^{2014}) + \varepsilon_{it},
 \end{aligned}
 \tag{1}$$

---

<sup>6</sup>Di Porto et al. (2025) consider as controls co-workers below the threshold and find that a non-negligible fraction left PA after the cap. In our interpretation, these managers are to be considered “later on” treated and, therefore, are inappropriate as controls.

where  $y_{it}$  is the log of annual salary for individual  $i$  in year  $t$ ,  $\alpha_i$  are individual fixed effects, and  $\delta_t$  are year fixed effects. The indicator  $\text{Eligible}_i^j$  equals one if individual  $i$  meets the eligibility criteria for policy  $j \in \{2011, 2014\}$ . The term  $\mathbb{1}\{\text{Year}_t = \tau\}$  is a year dummy equal to one when  $t = \tau$ , with 2010 as the omitted base year. The coefficients  $\beta_\tau^{2011}$  and  $\beta_\tau^{2014}$  capture the dynamic effect of eligibility for each policy relative to 2010.

These estimates should be interpreted as intention-to-treat (ITT) effects. Eligibility is determined by pre-policy salary thresholds and is therefore exogenous to policy implementation. Not all eligible individuals necessarily crossed the salary cap after the policy, so the coefficients measure the average impact of eligibility rather than the effect of actual treatment. This approach avoids bias from endogenous eligibility and remains valid under the parallel trends assumption. While ITT effects may be smaller than treatment-on-the-treated effects due to incomplete compliance, they are policy-relevant because they reflect the effect of eligibility under real-world conditions.

## 5 Empirical results

Figure 1 shows the effects of the two salary-cap policies introduced on log-salaries. Dotted vertical lines highlight the timing of each policy’s introduction. The figure shows that the parallel trends assumption holds before the policies were introduced and that the 2011 policy had the largest effect, with an immediate reduction in the average salary of eligible workers of over 16%, reaching 25% by the end of the period. The 2014 estimate also shows a significant effect on eligible individuals, but it reduces the average salary by less than 10%.

Because our research design uses an intention-to-treat (ITT) approach, it measures the impact of introducing a salary cap, regardless of whether it is actually affecting all eligible workers or only a part of them. Hence, the ITT effect provides the policy’s overall impact, as it includes both compliers and non-compliers.

Given our results, we compute a back-of-the-envelope estimate of the wage bill savings generated by the 2011 salary cap. In the EC data, we count 21,912 individuals earning more than €100,000 (at 2005 prices) in 2010, of whom 20% are affected by

the salary cap (Table 5). We inflate this count to reflect the random selection rate (48 days of birth over the year), and multiply the resulting figure by the average share of the eligible population (0.036; see Table 4) to obtain an estimate of the number of eligible workers. This is then multiplied by 0.461, the estimated likelihood of being affected by the cap, and by the mechanical salary drop (€134,560, see Section 6.1). This yields total gross salary savings of approximately €74.5 million per year, which declines to about €42.5 million after accounting for the applicable 43% income tax rate.<sup>7</sup>

## 5.1 Unintended effects

Besides moderating public expenditures for salaries in the PA, did salary cap policies produce any unintended or undeclared effects? Did they introduce an incentive for eligible workers to switch to a private-sector employer or retire earlier?

To answer the first question, we classify managers as switchers if they are public employees in year  $t$  and private employees in year  $t+1$ . This variable is a dichotomous variable equal to 1 if an individual moved to the private sector and is defined for the period 2005-2018, as our sample is limited to the 2005-2019 period, which is why the sample size is about 6,000 observations smaller (Table 6, top panel). Unconditional descriptive statistics also show that mobility out of the public sector is limited, on average affecting just 2.2% of the overall sample, reaching 3.7% for those eligible to the 2011 policy and 8.1% for those eligible only to the 2014 policy. The probability of switching to the private sector is lower (1.9%) among PAs' employees who are not eligible for salary caps.

To answer the second question and identify workers who choose to retire, we use an additional INPS archive, which records all pension payments made by INPS (*Archivio Pensioni*) to individuals. The pension archive was matched to the EC dataset using the unique individual identifier code. Pensions in Italy are often paid

---

<sup>7</sup>We are not aware of analogous calculations performed by others. Alternatively, one could adopt the approach typically used by the Italian Supreme Audit Institution responsible for auditing public expenditure (*Corte dei Conti*), although their assessments generally do not incorporate behavioural or dynamic responses and instead approximate the impact by multiplying the number of affected individuals in recent years by the mechanical gross salary loss.

without reference to a retirement age, based instead on a minimum number of years of contributions, in some cases allowing pension recipients to continue working even after receiving their first pension payment. This implies that some workers could be recorded as pensioners even while still working and earning a salary higher than their pension income. Hence, we classify an individual as a pensioner from the first year in which their pension income exceeds their public-sector employment income. Because the EC archive is not sufficiently complete to study retirement decisions, as it starts in 2005, when the oldest individuals were aged 55, we limit our analysis to individuals aged 55 or older.<sup>8</sup> Our sample including pensioners covers the period 2006–2019 and is slightly larger (96,623 observations; Table 6, bottom panel) than the employment income sample (92,673 observations; Table 4), because individuals who stop working exit the latter dataset, while they remain in the former as pensioners. The unconditional probability of being a pensioner in the whole sample in our 1950–1970 cohort of workers aged 55 or more is 1.4%, reaching 3.7% among those eligible to the 2011 policy and 8.2% for those eligible to the 2014 policy only.

The empirical model is as in (1), except that the dependent variable is replaced with either our dichotomous variable identifying switchers or our dichotomous variable identifying pensioners.

Figure 2 shows that the 2011 reform had a significant effect on increasing the likelihood of public employees with the highest salaries to switch to the private sector. Five years after the reform was introduced, the probability of leaving the public sector for the private sector was 5% higher and increased further in 2018. In contrast, the 2014 reform is estimated to have no significant effect on the likelihood that top PA employees exited the public sector.

Figure 3 confirms that the 2011 reform had a larger impact than the 2014 reform, also in the probability of early retirement. It increased by more than 5% in a few

---

<sup>8</sup>According to the pension reform introduced in 1995, a defined-benefit plan was replaced by a defined-contribution system, with a phase-out period. In addition, the retirement age was gradually raised to 67, unifying men’s and women’s exit ages. Early retirement was possible for workers with a minimum number of years of contributions (between 30 and 40 years). Since the old system was more generous than the new one, there was an incentive to retire as early as possible after the reform. So it was not impossible to observe early retirees at the beginning of our sample period, but most managers could retire effectively after age 60, i.e., after 2010.

years, reaching nearly 15% in 2019. The reduced incentive to keep working due to the salary cap was likely reinforced by the increasing likelihood of reaching the minimum contribution period required to claim a pension under the more generous earnings-related principle. Although the effect of the 2014 reform also shows a positive impact on the probability of retiring, it was always lower than that of the previous reform and was mostly not statistically significant.

These figures show that salary cap policies have indirect effects, prompting behavioural reactions by affected individuals. If switchers were less productive than stayers, the reforms could, in addition to saving on salaries, increase the average level of productivity in the public administration, or vice versa. Similarly, if earlier retirees were less productive than stayers, the salary cap reforms could increase average productivity while reducing the average age of the workforce. However, savings from paying salaries for shorter periods to top employees should be weighed against a larger expenditure on their pensions over longer periods, assuming equal life expectancy, which might yield an unclear public finance gain, as retired workers need to be replaced to some extent.

To investigate this issue, we performed the following simple exercise. We run an OLS regression of log employment income on a full set of year dummies for the pre-reform period (2005-2010) and an individual fixed effect. The estimated individual fixed effect, controlling for time trends, can be interpreted as a measure of time-invariant individual productivity.

Figure 4 contrasts the distribution of individual productivity for stayers and switchers (left panel) and for stayers and retirees (right panel). It shows that workers who switch from the public to the private sector exhibit greater variability but also a higher mean. In other words, it shows that the behavioural effect of a salary cap is heterogeneous, with a stronger reaction among workers with a higher proxy for individual productivity. The exit of most productive workers from the PA is likely to harm the PA's efficiency, at least in the short run, before they are possibly, if ever, replaced by workers at least as productive. In a similar vein, we observe that managers who retire after the cap is introduced are more productive than stayers.

## 5.2 Robustness analysis

In Section 4, we defined eligible individuals for the 2011 policy as those who received salary above € 225,000 for at least one year in the period 2005-2010 and worked for a public sector institution where the salary cap applies, and eligible individuals for the 2014 policy as those who received salary above € 175,000 for at least one year in the period 2005-2013, work for a public sector institution where the salary cap applies, and were not eligible for the 2011 policy, as defined above. Here, we adjust both salary thresholds for eligibility upward and downward by € 25,000. By moving down the eligibility thresholds, we increase the eligibility sample size and consequently also the precision of the policy effect estimate, confirming the main results presented in Section 5, except for the probability of retiring after the 2014, which is now significantly positive, and only slightly smaller than the effect of the 2011 reform. Increasing the eligibility threshold by € 25,000 reduces the sample size of the eligibility group, increasing the variability of the estimate. No notable difference with respect to what is presented in Section 5 emerges: the effects of the 2011 reform are always significant, and those of the 2014 reform are smaller and not statistically significant (Figures 5, 6 and 7).

By moving down and up both the eligibility thresholds, and following the same procedure to define the individual worker's productivity as explained in Section 5.1, we also confirm that the proxy for individual productivity who switch to the private sector or retire, following the salary cap reforms is larger than that of stayers (Figures 8 and 9).

## 6 Implications of the salary cap policy: elasticities, and welfare analysis

We adopt a standard framework (Kaplow, 2024) in which utility depends on disposable income and labour effort:

$$u = u(c, l), \quad u_c > 0, \quad u_l < 0.$$

Let  $l$  denote hours of work in the public sector and let  $y = wl$  denote gross income from public-sector employment, where  $w$  is the wage rate. Disposable income is given by a linearized tax system,

$$c = zy + V,$$

where  $z \in (0, 1)$  is the marginal net-of-tax rate and  $V$  is the virtual income, i.e. the intercept of the linearized budget constraint obtained by approximating a nonlinear tax schedule at a given income level.<sup>9</sup> The behavioural choice variable on the intensive margin is public-sector earnings  $y$ , with policy changes operating through the marginal incentive  $z$ . Since  $y = wl$ , the choice can equivalently be represented as hours of work or taxable income, as is standard in the taxable-income literature.

The individual chooses labour  $l$  and a discrete status  $d \in \{\text{stay, switch, retire}\}$ :

$$(l, d) \in \arg \max \{U_{\text{stay}}, U_{\text{switch}}, U_{\text{retire}}\}.$$

Utilities are defined as  $U_{\text{stay}}(c, l) = u(zy + V, l)$ ,  $U_{\text{switch}}(c, l) = u(c_{\text{switch}}, 0) - \kappa_{\text{switch}}$ ,  $U_{\text{retire}}(c, l) = u(c_{\text{retire}}, 0) - \kappa_{\text{retire}}$ . When  $d \in \{\text{switch, retire}\}$ , the individual supplies no public-sector labor, but consumes an exogenous level of income  $c_{\text{switch}}$  or  $c_{\text{retire}}$  reflecting earnings or pensions outside public employment. The parameters  $\kappa_{\text{switch}}, \kappa_{\text{retire}} \geq 0$  capture non-pecuniary exit costs associated with switching jobs or retiring.<sup>10</sup>

Before the cap binds, the intensive-margin optimum ( $y_0 < \bar{y}$ ) for stayers satisfies

$$u_c \frac{\partial c}{\partial l} + u_l = 0, \quad \frac{\partial c}{\partial l} = \frac{\partial c}{\partial y} \frac{\partial y}{\partial l} = zw,$$

---

<sup>9</sup>This ensures that the linear approximation preserves both marginal incentives and disposable income at that point. Since we focus on salaries well within the top income bracket, the budget constraint is linear over the relevant range, and the virtual-income term is constant. The virtual income is computed as the difference of  $(1 - z)y$  and the actual tax paid with the progressive tax system, with  $y$  equal to the lower bound of the top income tax bracket.

<sup>10</sup>Public-sector jobs offer limited bargaining scope and virtually no tax-avoidance channels compared with private top-income settings. Public managers could, in principle, bypass the cap through private consultancy work, which is not restricted by the cap. These occasional earnings are recorded in the EC archives under *gestione separata*. Using EC data, we find no significant change in the incidence of such earnings after the introduction of the 2011 and 2014 caps.

so the first-order condition is

$$u_c(zw) + u_l = 0.$$

A hard cap on gross income  $\bar{y}$  implies that  $y(l) = \min\{wl, \bar{y}\}$ . Hence, disposable income, after a salary cap is introduced (time  $t = 1$ ), becomes

$$c_1 = z \min\{wl, \bar{y}\} + V.$$

For  $wl < \bar{y}$ , the budget line has slope  $zw > 0$ , whereas for  $wl \geq \bar{y}$  it becomes flat, meaning a 100% marginal tax rate above  $\bar{y}$ , so any additional labour beyond  $y = \bar{y}$  yields no further disposable income.

Since utility is strictly increasing in  $c$  and strictly decreasing in  $l$ , any choice  $l > \bar{y}/w$  is strictly dominated. Hence, stayers whose unconstrained optimum satisfies  $y_0 > \bar{y}$  optimally locate at the kink point  $l_1 = \bar{y}/w$ .

Large and discrete changes in the marginal net-of-tax rate—such as those induced by salary caps or notches—render the local derivative  $d \log y / d \log z$  ill-defined, motivating a finite-change interpretation of behavioural responses (Saez et al., 2012; Kleven and Waseem, 2013). In our empirical work, however, we follow standard top-bracket applications and report a normalized taxable-income elasticity, relating percentage changes in gross income to percentage changes in the marginal net-of-tax rate computed at the pre-reform level:

$$\varepsilon = \frac{\Delta y / y_0}{\Delta z / z_0}.$$

Normalizing by the pre-reform marginal incentive ensures consistency with sufficient-statistics formulas and maintains comparability with existing top-bracket applications (Saez, 2010; Piketty et al., 2014).

In our empirical context, at high earnings, the pre-cap marginal tax rate is 43%, hence  $z_0 = 0.57$ . Under the cap, the marginal net-of-tax rate above  $\bar{w}$  collapses to 0, so  $\Delta z / z_0 = (0 - 0.57) / 0.57 = -1$ . To estimate the average effect on eligible workers after the introduction of the salary cap, we run a difference-in-differences

regression including a post-2011 indicator, an eligibility dummy for individuals with an annual earnings above €225,000, their interactions, and individual fixed effects. The coefficient on the interaction term yields an ITT estimate of  $-0.214$ , hence the average intensive (stayers’) taxable income elasticity reported in top-bracket form is

$$\varepsilon \approx \frac{\Delta y/y_0}{\Delta z/z_0} = 0.214. \quad (2)$$

For extensive margins, let  $P_{\text{switch}}$  and  $P_{\text{retire}}$  be the transition probabilities. Regressing each on the post-2011 dummy, a difference-in-differences regression including a post-2011 indicator, an eligibility dummy for individuals with an annual earning above €225,000, their interactions, and individual fixed effects gives  $\Delta P_{\text{switch}} = 0.031$  and  $\Delta P_{\text{retire}} = 0.074$ . Using  $\Delta z/z_0 = -1$ , and  $P_{\text{switch},0} = 0.022$  and  $P_{\text{retire},0} = 0.014$ , the conventional top-bracket elasticities are interpreted as finite-change elasticities and reported in absolute value:<sup>11</sup>

$$\varepsilon^{\text{switch}} = \left| \frac{\Delta P_{\text{switch}}/P_{\text{switch},0}}{\Delta z/z_0} \right| = 1.41, \quad \varepsilon^{\text{retire}} = \left| \frac{\Delta P_{\text{retire}}/P_{\text{retire},0}}{\Delta z/z_0} \right| = 5.29.$$

The extensive elasticities are substantially larger than the intensive elasticity. While [Piketty et al. \(2014\)](#) focus on avoidance and bargaining rather than switching or retirement, their “three elasticities” framework emphasizes that top-income responses need not operate only through the intensive labour-supply margin. Our findings echo this insight: the adjustment to the salary-cap reform also occurs through switching and retiring, not only through earnings adjustments among stayers. In our setting, switching and retirement constitute the main channel through which the policy operates.

---

<sup>11</sup>These extensive-margin elasticities are conceptually analogous to the migration elasticities studied by [Kleven et al. \(2013\)](#) and to the finite-change responses analyzed by [Kleven and Waseem \(2013\)](#). In those settings, large and discrete tax changes induce reallocations across discrete alternatives rather than marginal adjustments to a continuous choice. The elasticities reported here, therefore, normalize proportional changes in transition probabilities by the proportional change in the net-of-tax rate at the pre-reform level, and should be interpreted as finite-change responses rather than local Marshallian elasticities.

Although extensive-margin elasticities exceeding one may appear unusual when compared to standard intensive-margin elasticities, they are entirely consistent with the theory and evidence on high-income labour supply under large and discrete policy changes. First, these elasticities measure *proportional* changes in transition probabilities, so whenever the baseline probability  $P_{\text{switch},0}$  or  $P_{\text{retire},0}$  is small, as is typically the case for top public-sector managers, even moderate absolute changes imply large percentage responses. Second, in environments with notches, salary caps, or confiscatory marginal rates, as when the net-of-tax rate collapses from  $z_0$  to  $z_1 = 0$ , the standard envelope-theorem logic behind smooth intensive adjustments no longer applies. The individual is no longer marginally adjusting an interior choice but instead re-optimising over discrete alternatives, such as remaining in the public sector, switching to the private sector, or retiring. In such settings, extensive-margin elasticities can naturally exceed one because they reflect shifts across distinct utility bundles rather than incremental changes in hours or effort.

Third, the empirical public-finance literature routinely documents extensive responses of this magnitude at the top of the income distribution. Large extensive-margin responses to discrete shifts in lifetime resources are also documented in settings unrelated to taxation, such as wealth shocks: [Brülhart et al. \(2025\)](#) show that inheritance-induced increases in net wealth lead to sizeable reductions in labour supply and markedly higher retirement hazards among older high-income individuals. Large mobility elasticities are reported in studies of high-income migration and preferential tax regimes ([Kleven et al., 2013](#)), and similarly large, discrete adjustments arise in notch settings where taxpayers respond non-locally to fiscal discontinuities ([Kleven and Waseem, 2013](#)). More broadly, analyses of top earners show that behavioural responses often operate through non-real-supply margins, such as bargaining over rents, income shifting, or broader mobility choices, rather than smooth adjustments in effort, as emphasised by [Piketty et al. \(2014\)](#). Although our setting features different discrete alternatives (sector switching or retirement), the underlying mechanism is similar: when policy introduces a sharp discontinuity in the budget set, adjustments occur primarily through changes in employment or occupational status rather than through incremental variation in hours or earnings. Similar patterns

also arise in macroeconomic models with indivisible labour, where extensive-margin elasticities calibrated to match aggregate fluctuations are much larger than intensive ones (Rogerson, 1988; Hansen, 1985; Chetty et al., 2011). In all these cases, values greater than one are not anomalies but the expected outcome of a discrete choice under large fiscal discontinuities. Our estimated extensive elasticities, therefore, reflect a well-understood mechanism: when a policy creates a sharp drop in the net-of-tax return (here, a 100% marginal tax rate above the salary cap) the dominant behavioural response is on the extensive margin, and its elasticity can readily exceed unity.

## 6.1 Welfare accounting

A hard cap on gross income  $\bar{y}$  truncates the budget set and makes disposable income flat for  $y > \bar{y}$ . Among stayers, the cap is binding only for individuals whose unconstrained pre-cap optimum satisfies  $y_0 > \bar{y}$ . For this group, the only feasible choice in the flat region that is not strictly dominated is the kink point, since any  $l > \bar{y}/w$  yields the same disposable income but strictly higher effort cost. Thus their post-cap optimal labour supply is the corner solution  $l_1 = \bar{y}/w$ , gross income is  $y_1 = \bar{y}$  and post-cap disposable income is

$$\bar{c} = z_0 \bar{y} + V.$$

For stayers with  $y_0 \leq \bar{y}$ , the cap does not bind and  $l_1 = l_0$ .

For notational convenience, define net-of-intercept disposable income

$$\tilde{c} \equiv c - V = z_0 y = z_0 w l.$$

Accordingly,  $\tilde{c}_0 = z_0 w l_0 = z_0 y_0$  and  $\tilde{c} = z_0 \bar{y}$ . The welfare loss for a stayer is

$$\Delta W^{\text{intensive}} \equiv u(c_0, l_0) - u(\bar{c}, l_1). \quad (3)$$

Let  $\Delta c \equiv \bar{c} - c_0$  and  $\Delta l \equiv l_1 - l_0$ . Along the pre-cap budget line  $c = z_0 y + V = z_0 w l + V$ , we have  $\Delta c = z_0 w \Delta l$ .

A second-order expansion of  $u(c, l)$  around  $(c_0, l_0)$  gives

$$\Delta W^{\text{intensive}} \approx \frac{1}{2} \left[ -u_{cc}(c_0, l_0) (\Delta c)^2 - 2u_{cl}(c_0, l_0) \Delta c \Delta l - u_{ll}(c_0, l_0) (\Delta l)^2 \right], \quad (4)$$

with the first-order term vanishing at the optimum.

Substituting  $\Delta l = \Delta c / (z_0 w)$  yields

$$\Delta W^{\text{intensive}} \approx \frac{1}{2} (\Delta c)^2 \left[ -u_{cc} - \frac{2u_{cl}}{z_0 w} - \frac{u_{ll}}{(z_0 w)^2} \right]_{(c_0, l_0)}. \quad (5)$$

To express the welfare loss in euros of disposable income, we adopt the standard quasi-linear specification  $u(c, l) = c - C(l)$ , where  $C(l) = \frac{1}{2} a l^2$  (e.g. [Saez et al., 2012](#); [Kleven, 2016](#)). Under quasi-linearity, welfare is measured in euros of disposable income, so the difference between pre-cap utility and utility under the cap corresponds to the area between marginal benefit and marginal cost. The unconstrained optimum  $l_0$  satisfies  $a l_0 = z_0 w$ , and with  $\tilde{c}_0 = z_0 w l_0$  this implies  $a = (z_0 w)^2 / \tilde{c}_0$ .

Evaluating welfare at the unconstrained optimum  $y_0$  and at the capped choice  $\bar{y}$  gives the Harberger deadweight-loss triangle (Figure 10). The welfare difference can be written as the integral  $\int_{\bar{y}}^{y_0} [z_0 - \frac{a}{w^2} y] dy$ , i.e. the wedge between marginal benefit and marginal cost. Using a second-order Taylor expansion of utility around the optimum, the welfare loss is approximated by

$$\begin{aligned} \Delta W^{\text{intensive}} &= W_0 - W(\bar{y}) = \int_{\bar{y}}^{y_0} \left[ z_0 - \frac{a}{w^2} y \right] dy \\ &\approx \frac{1}{2} (\Delta c)^2 \left[ \frac{a}{(z_0 w)^2} \right] = \frac{(\tilde{c}_0 - \bar{c})^2}{2 \tilde{c}_0} = \frac{(c_0 - \bar{c})^2}{2(c_0 - V)}. \end{aligned} \quad (6)$$

Thus, the intercept  $V$  cancels from the mechanical income change in the numerator and enters only through the denominator, via  $\tilde{c}_0 = c_0 - V$ . Under these assumptions,  $\Delta W^{\text{intensive}}$  is the welfare loss associated with the cap over  $[\bar{y}, y_0]$ .

Given an estimated average ITT of  $-0.214$  and a 46.1% estimated likelihood of being affected by the cap, the implied LATE is  $-0.464$ . This corresponds to a pre-reform salary  $y_0$  that was 46.4% higher than the cap  $\bar{y}$ . Given  $\bar{y} = \text{€}290,000$ , we

have  $y_0 = 1.464\bar{y} = \text{€}424,560$ . Using  $V = \text{€}6,830$  and  $z_0 = 0.57$ , we obtain:

$$\bar{c} = z_0\bar{y} + V = 0.57 \times \text{€}290,000 + \text{€}6,830 \approx \text{€}172,000,$$

$$c_0 = z_0y_0 + V = 0.57 \times \text{€}424,560 + \text{€}6,830 \approx \text{€}249,000.$$

Hence,

$$c_0 - \bar{c} \approx \text{€}77,000,$$

and

$$c_0 - V = z_0y_0 = \tilde{c}_0 \approx \text{€}242,000.$$

Applying (6), the intensive-margin welfare loss for the average affected worker is:

$$\Delta W^{\text{intensive}} = \frac{(77,000)^2}{2 \times 242,000} \approx 12,250.$$

A simple back-of-the-envelope calculation clarifies the magnitude of the implied effort response. Before the salary cap, average net-of-tax earnings were approximately  $\text{€}249,000$ . Public-sector contractual working time is 173 hours per month (over 12 months), corresponding to 40 hours per week, although it is reasonable to assume that top managers work around 10-15 hours more per week. A reduction in effort costs of  $\text{€}64,500$  corresponds to a decrease of roughly 650–800 annual working hours, depending on the assumed net-of-tax hourly value of time  $z_0w$ . This suggests that top managers reduced their optimal working hours essentially down to the minimum contractual level.

Let  $U_0$  denote the pre-policy *private surplus from labour* (the area under marginal surplus from 0 to  $l_0$ ). With  $u(c, l) = c - C(l)$  and  $C(l) = \frac{1}{2}al^2$ ,

$$U_0 = \int_0^{l_0} (z_0w - C'(l)) dl = \frac{(z_0w)^2}{2a} = \frac{\tilde{c}_0}{2} = \frac{c_0 - V}{2}.$$

Using the observed changes in exit probabilities, we define:

$$\Delta W^{\text{extensive}} = \Delta P_{\text{switch}}(U_0 - U_{\text{switch}}) + \Delta P_{\text{retire}}(U_0 - U_{\text{retire}}). \quad (7)$$

We normalize the outside-option utilities to zero,  $U_{\text{switch}} = U_{\text{retire}} = 0$ .<sup>12</sup> In our setup, these utilities differ from zero only because they include non-pecuniary costs of switching jobs or retiring (collected in the  $\kappa$ -terms). Since the extensive-margin welfare loss depends solely on the differences  $(U_0 - U_{\text{switch}})$  and  $(U_0 - U_{\text{retire}})$ , these constants can be absorbed into the utility baseline without affecting behavior or welfare comparisons. Under this normalization, an induced exit destroys exactly the worker's private surplus from employment,  $U_0$ .

$$\Delta W^{\text{extensive}} = (\Delta P_{\text{switch}} + \Delta P_{\text{retire}}) U_0 = (\Delta P_{\text{switch}} + \Delta P_{\text{retire}}) \frac{c_0 - V}{2}. \quad (8)$$

With  $c_0 - V \approx \text{€}242,000$  we have

$$U_0 = \frac{c_0 - V}{2} \approx \text{€}121,000.$$

Using  $\Delta P_{\text{switch}} = 0.031$  and  $\Delta P_{\text{retire}} = 0.074$ ,

$$\Delta W^{\text{extensive}} = 0.105 \times \text{€}121,000 \approx \text{€}12,700.$$

Summing intensive and extensive components,

$$\Delta W^{\text{total}} = \Delta W^{\text{intensive}} + \Delta W^{\text{extensive}} \approx \text{€}12,250 + \text{€}12,700 \approx \text{€}24,950$$

Combining the two margins, the earnings cap generates an overall welfare loss that

---

<sup>12</sup>This normalisation implies that we treat the exit as a total loss of the private surplus from public-sector employment. A more general specification would define  $\Delta W^{\text{extensive}}$  as the difference between the public-sector surplus and the (potentially taxable) surplus in the outside option. Our calculation, therefore, provides a benchmark for the private utility destruction induced by the cap, though the net social welfare loss would be lower if the outside options generate significant fiscal externalities, such as taxes on private-sector earnings or pensions.

remains small relative to both the level of pre-cap net earnings and the mechanical income reduction. On the intensive margin, the welfare loss from truncating earnings is approximately €12,250, corresponding to about 5.0% of pre-cap net-of-tax salary ( $\tilde{c}_0 \approx \text{€}242,000$ ). This reflects the fact that only the local curvature of preferences around the pre-cap optimum matters for welfare. The extensive margin contributes a similar amount, about €12,700, or roughly 5.2% of pre-cap net earnings, driven by increased switching and retirement. Taken together, the total welfare loss is approximately €24,950 per affected worker, corresponding to about 10.3% of pre-cap net-of-tax salary. These magnitudes are consistent with the second-order nature of behavioral responses to high-income kinks and are robust under alternative convex specifications of effort costs.

## 7 Conclusions, discussion and policy implications

This paper has examined the effects of Italy’s salary caps for top public managers using linked administrative records and a difference-in-differences event study design that exploits institutional exemptions as a comparison group. The 2011 reform—by imposing a binding ceiling and raising the marginal tax rate from 43% to 100% above the threshold—generated immediate and persistent compression of top public-sector earnings. By contrast, the 2014 tightening had comparatively smaller effects, underscoring the primacy of the first, more salient intervention. Our empirical strategy treats pre-policy threshold exceedance as an eligibility instrument and estimates intention-to-treat effects at policy onsets, thereby capturing policy-relevant responses under real-world compliance.

Three substantive findings emerge. First, salary caps achieved their primary fiscal objective. Affected individuals experienced an average annual reduction in gross earnings of more than €130,000. When scaled to the full population, these reductions translate into aggregate gross savings of more than 74 million per year. Because forgone earnings would otherwise have been taxed at 43%, the net-of-tax savings remain sizable at roughly €42 million annually. These figures highlight the fiscal significance of the reform and the importance of distinguishing between gross and net savings.

Second, intensive-margin responses are modest. Despite the sharp drop in the marginal net-of-tax rate from  $z_0 = 0.57$  to zero, taxable-income responses at the top are limited. We estimate an intensive elasticity of approximately 0.214 and compute an associated welfare loss of about 12,250 per capita-affected individual under a standard sufficient-statistics calibration. In this framework, efficiency costs reflect the reduction in private surplus net of saved effort costs, rather than the full mechanical income loss, so deadweight losses are second-order in the change in earnings (Saez, 2010; Chetty, 2009). Consistent with theory, the presence of a salary cap does not imply large intensive distortions when effort is only weakly contractible.

Third, extensive-margin adjustments dominate both behavior and welfare. The 2011 cap significantly increased the probabilities of switching to private-sector employment and of early retirement.

Several welfare-relevant channels are unobserved in available administrative data, implying that our estimates likely understate true social costs. First, we cannot directly measure public-service quality. High-skill managers may generate social value that exceeds their private surplus, so cap-induced exits could degrade organizational performance (e.g. Fenizia, 2022; Limodio, 2021; Lazear et al., 2015). Second, replacement and training costs—hiring delays, onboarding, learning curves, and disruptions to team production—represent real but unmeasured efficiency losses. Third, distinguishing productive effort from rent extraction is challenging in public administration, where outputs lack market prices and many services (justice, health, defense, education) are monopolistic. Fourth, general-equilibrium spillovers may arise within agencies and across sectors, affecting morale, internal promotion, and private-sector absorption. Fifth, dynamic selection matters: the deterrence of future entrants into public management may have long-run effects that exceed short-run elasticities suggest (Chetty, 2012). Sixth, early retirements reshape pension liabilities and the timing of fiscal flows. Finally, while we find no evidence of increased consultancy or temporary-contract income following the cap, other forms of norm avoidance may exist—especially in settings where SOE remuneration or income reclassification is feasible (OECD, 2024; Huang and Villadsen, 2023). Taken together, these omitted or hard-to-measure margins indicate that our welfare estimates should be interpreted

as conservative lower bounds.

Two caveats qualify interpretation. First, our salary measures are derived from accrual-basis contribution records, whereas the cap is enforced on a cash-year basis; this blurs visible kinks at the threshold. Nevertheless, the sufficient-statistics approach relies on average responses rather than bunching and remains valid in this environment (Saez, 2010; Chetty, 2009). Second, we cannot observe manager-level output; hence, productivity selection is inferred from pre-reform individual fixed effects rather than value-added, which limits our ability to separate efficiency from rent extraction.

External validity requires care. Italy combines limited opportunities for avoidance, heterogeneous exemptions (including constitutional bodies and joint-stock SOEs), and fixed-term managerial appointments. In settings with broader avoidance or different exemption maps, the decomposition across effort, avoidance, bargaining, and participation margins may shift. Even so, a general lesson emerges: when salary caps effectively impose confiscatory marginal rates at the very top, intensive deadweight losses need not be large, but participation responses can generate substantial social costs through the exit or reallocation of high-value managers (Piketty et al., 2014).

Table 1: Number of public administrations (PAs) and of workers, with and without salary cap, by main categories of PAs, using EC and Uniemens datasets in year 2019.

	EC				Uniemens			
	Organisations		Workers*		Organisations		Workers	
	Salary cap	No salary cap	Salary cap	No salary cap	Salary cap	No salary cap	Salary cap	No salary cap
Italian Automobile Club	20	0	2,928	0	50	0	2,001	0
Tax agencies	4	0	55,298	0	4	0	9,001	0
Authorities	33	0	4,266	0	23	0	1,677	0
Chambers of Commerce	81	0	5,848	0	14	0	205	0
Municipalities	6,100	0	398,770	0	944	0	4,350	0
Consortia	387	0	13,246	0	292	0	3,583	0
Authorities, instit., agencies	557	0	88,847	0	292	0	13,990	0
Foundations	389	0	16,874	0	580	0	9,033	0
Ministries	15	0	1,329,490	0	5	0	1,106	0
Constitutional bodies	8	2	4,540	7,475	3	2	327	7,211
Social security	22	0	43,496	0	22	0	39,661	0
Professional orders	336	0	3,536	0	1,081	0	2,676	0
Provinces-Regions	126	0	116,610	0	35	0	1,847	0
Healthcare	214	0	672,832	0	123	0	2,802	0
Services to people	401	0	31,033	0	89	0	734	0
PLC, other companies	3,319	54	604,296	285,293	4,280	55	536,962	285,712
Sport, culture	108	0	13,733	0	37	0	761	0
Universities, Schools	588	0	188,089	0	60	0	2,219	0
Unclassified	26	0	6,441	0	45	0	2,095	0
<b>Total</b>	<b>12,734</b>	<b>56</b>	<b>3,600,170</b>	<b>292,768</b>	<b>7,979</b>	<b>57</b>	<b>635,030</b>	<b>292,923</b>

Notes: The categories of public administration (PA) have been classified by exact matching based on individual institutions' tax codes of external information with the Estratti Conto (EC) and Uniemens datasets.

\* To facilitate comparison with aggregate statistics, workers' figures obtained from the EC dataset are divided by 48/365, as the EC covers a random sample of the workforce consisting of individuals born on 48 dates in each calendar year.

Table 2: Some descriptive statistics, workers in PA and workers in the private sector, using the whole EC dataset for year 2019, all workers and workers in the 1950-1970 cohort.

Sector	Whole sample		1950-1970 birth cohort	
	PA (A)	Private (B)	PA (C)	Private (D)
<i>Panel a: Demographic characteristics</i>				
Female share	0.60	0.41	0.59	0.39
Foreign share	0.02	0.15	0.01	0.10
Age (mean)	49.82	41.55	56.97	55.14
<i>Panel b: Income variables (in current € )</i>				
Mean	31,957	21,957	35,321	27,183
St. dev.	21,615	25,864	23,652	28,958
10th percentile	12,296	2,471	18,118	4,465
50th percentile	29,444	19,284	31,544	22,904
90th percentile	50,255	42,188	54,515	50,535
99.5th percentile	123,611	113,070	139,306	162,427
N. obs.	510,640	1,994,360	302,143	631,747

Notes: Statistics computed using *Estratti Conto* (EC) dataset, in year 2019.

Table 3: The probability of being over the 2011 and 2014 salary cap threshold in 2010, for PA employees receiving different levels of salary in 2005.

		<b>Year 2010</b> over € 290,000
<b>Year 2005</b>	over € 200,000	37.9%
	over € 225,000	46.1%
	over € 250,000	50.5%
	over € 290,000	56.5%
		<b>Year 2010</b> over € 240,000
<b>Year 2005</b>	over € 150,000	30.8%
	over € 175,000	42.5%
	over € 200,000	52.4%
	over € 240,000	60.5%

Notes: The table presents the unconditional probability of a public-sector employee whose salary falls in a given income tax bracket in 2005 to be over the 2011 and 2014 salary cap threshold (respectively, €290,000 and €240,000) in 2010, before any salary cap policy was introduced.


Table 4: Descriptive statistics of main variables used in the regressions

<b>Variable</b>	<b>Obs</b>	<b>Mean</b>	<b>St.d.</b>	<b>Min</b>	<b>Max</b>
<i>Whole sample</i>					
Salary (log)	92,673	11.548	0.572	-4.605	15.284
Eligible 2011 policy	92,673	0.036	0.187	0.000	1.000
Eligible 2014 policy	92,673	0.081	0.273	0.000	1.000
<i>No salary cap</i>					
Salary (log)	13,131	11.662	0.676	2.565	14.254
<i>salary cap</i>					
Salary (log)	79,542	11.529	0.550	-4.605	15.284
Eligible 2011 policy	79,542	0.042	0.201	0.000	1.000
Eligible 2014 policy	79,542	0.094	0.292	0.000	1.000

---

Notes: The sample consists of individuals working in the public sector, born between 1950 and 1970, who received a yearly salary above € 100,000 at least once during the period 2005–2019 (in 2005 euros, CPI-adjusted). Only individuals working for the public sector, born between 1950 and 1970, are considered. The sample includes all individuals who, at least once during 2005-2019, received an annual salary above 100,000 (in 2005 euros, CPI-adjusted). Eligible individuals for the 2011 policy are those who earned salary above € 225,000 for at least one year between 2005 and 2010 and who work for a public-sector institution subject to the salary cap. Eligible individuals for the 2014 policy are selected as those who received salary above € 175,000 for at least one year in the period 2005-2013, work for a public sector institution where the salary cap applies, and were not eligible for the 2011 policy, as defined above.

Figure 1: Effects of the 2011 and 2014 policies on log-salaries.



Notes: See notes in Table 4.

Table 5: All workers, with salary above the € 100.000 of year 2005 (CPI adjusted).

Year	N. obs	% in PA, All	% in PA, with salary cap
2005	23,142	0.22	0.18
2006	23,171	0.22	0.18
2007	23,002	0.24	0.19
2008	22,800	0.24	0.19
2009	22,391	0.24	0.19
2010	21,912	0.25	0.20
2011	21,429	0.26	0.20
2012	20,852	0.26	0.20
2013	20,221	0.26	0.20
2014	19,616	0.26	0.20
2015	19,171	0.26	0.20
2016	18,526	0.26	0.20
2017	17,785	0.25	0.19
2018	16,916	0.24	0.18
2019	15,970	0.23	0.17


Notes: This table shows the yearly frequency of all workers in the EC dataset with income above €100,000 (at 2005 prices, CPI-adjusted), the share of those working in the public administration (PA), and in the PA with salary cap.

Table 6: Descriptive statistics of main variables used in the regressions

<b>Variable</b>	<b>Obs</b>	<b>Mean</b>	<b>St.d.</b>	<b>Min</b>	<b>Max</b>
<i>Whole sample</i>					
Prob. of switching to the private sector	86,348	0.022	0.146	0.000	1.000
Eligible 2011 policy	86,348	0.037	0.189	0.000	1.000
Eligible 2014 policy	86,348	0.081	0.274	0.000	1.000
<i>No salary cap</i>					
Prob. of switching to the private sector	12,091	0.019	0.135	0.000	1.000
<i>salary cap</i>					
Prob. of switching to the private sector	74,257	0.022	0.148	0.000	1.000
Eligible 2011 policy	74,257	0.043	0.203	0.000	1.000
Eligible 2014 policy	74,257	0.095	0.293	0.000	1.000
<i>Whole sample</i>					
Probability of being a pensioner	93,623	0.014	0.119	0.000	1.000
Eligible 2011 policy	93,623	0.037	0.189	0.000	1.000
Eligible 2014 policy	93,623	0.082	0.274	0.000	1.000
<i>No salary cap</i>					
Probability of being a pensioner	13,215	0.016	0.124	0.000	1.000
<i>salary cap</i>					
Probability of being a pensioner	80,408	0.014	0.119	0.000	1.000
Eligible 2011 policy	80,408	0.043	0.203	0.000	1.000
Eligible 2014 policy	80,408	0.095	0.293	0.000	1.000


Notes: The sample consists of individuals working in the public sector, born between 1950 and 1970, who received a yearly salary above € 100,000 at least once during the period 2005–2019 (in 2005 euros, CPI-adjusted). An individual is defined as a switcher if employed in the public sector in year  $t$  and observed in the private sector in year  $t + 1$ . An individual is classified as a pensioner starting in the first year in which pension income exceeds employment income; the sample is further restricted to individuals aged 55 or over. Eligible individuals for the 2011 policy are those who received a salary above € 225,000 for at least one year during 2005–2010 and worked for a public-sector institution subject to the salary cap. Eligible individuals for the 2014 policy are those who received a salary above € 175,000 for at least one year during 2005–2013, worked for a public-sector institution subject to the salary cap, and were not eligible for the 2011 policy, as defined above.

Figure 2: Effects of the 2011 and 2014 policies on the probability of switching to the private sector.



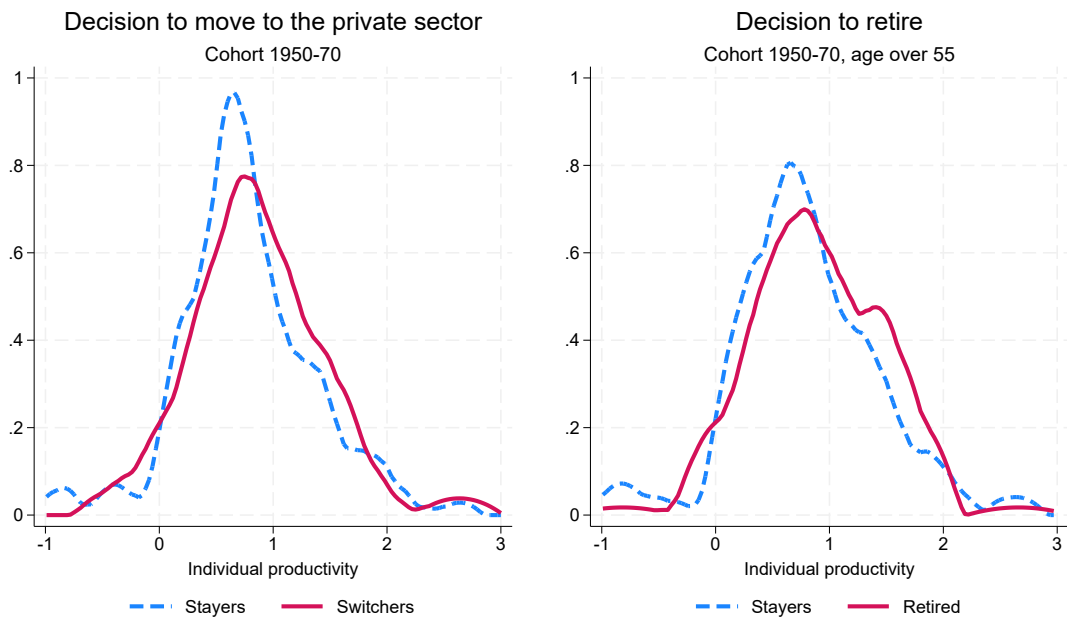
Notes: See notes in Table 6.

Figure 3: Effects of the 2011 and 2014 policies on the probability of retiring.



Notes: See notes in Table 6.

Figure 4: Density functions of individual productivity for civil servants who switch to the private sector/retire vs those who stay.




See notes in Table 6.

Table 7: Some descriptive statistics, workers in PA and workers in the private sector, using the Uniemens dataset for year 2019, all workers and workers in the 1950-1970 cohort.


Sector	Whole sample		1950-1970 birth cohort	
	PA	Private	PA	Private
<i>Demographic characteristics</i>				
Female share	0.38	0.42	0.35	0.39
Foreign share	0.02	0.16	0.02	0.11
Age (mean)	47.53	41.11	56.56	55.17
<i>salary variables</i>				
Mean	36,162	20,905	41,394	25,941
St. dev.	28,978	21,595	35,705	29,327
10th percentile	11,202	2,603	16,483	4,418
50th percentile	32,961	18,610	36,858	22,033
90th percentile	57,781	38,811	64,814	47,937
99.5th percentile	162,952	103,928	201,852	154,339
N. obs.	979,984	14,655,685	491,855	4,437,926

Notes: Statistics computed using Uniemens dataset, in year 2019.

Figure 5: Robustness for (log-) employment income




(a) Shifting the eligibility threshold down by € 25,000.




(b) Shifting the eligibility threshold up by € 25,000.

Notes as in Figure 1. In Panel 5a, eligible individuals for the 2011 policy are selected as those who received salary above € 200,000 for at least one year in the period 2005-2010 and work for a public sector institution where the salary cap applies. Eligible individuals for the 2014 policy are selected as those who received salary above € 150,000 for at least one year in the period 2005-2013, work for a public sector institution where the salary cap applies, and were not eligible for the 2011 policy, as defined above. In Panel 5b, salary eligibility thresholds are respectively moved to € 250,000 and € 200,000.

Figure 6: Robustness for probability of switching to the private sector




(a) Shifting the eligibility threshold down by € 25,000.




(b) Shifting the eligibility threshold up by € 25,000.

Notes as in Figure 2. In Panel 6a, eligible individuals for the 2011 policy are selected as those who received salary above € 200,000 for at least one year in the period 2005-2010 and work for a public sector institution where the salary cap applies. Eligible individuals for the 2014 policy are selected as those who received salary above € 150,000 for at least one year in the period 2005-2013, work for a public sector institution where the salary cap applies, and were not eligible for the 2011 policy, as defined above. In Panel 6b, salary eligibility thresholds are respectively moved to € 250,000 and € 200,000.

Figure 7: Robustness for probability of retiring




(a) Shifting the eligibility threshold down by € 25,000.



(b) Shifting the eligibility threshold up by € 25,000.


Notes as in Figure 3. In Panel 7a, eligible individuals for the 2011 policy are selected as those who received salary above € 200,000 for at least one year in the period 2005-2010 and work for a public sector institution where the salary cap applies. Eligible individuals for the 2014 policy are selected as those who received salary above € 150,000 for at least one year in the period 2005-2013, work for a public sector institution where the salary cap applies, and were not eligible for the 2011 policy, as defined above. In Panel 7b, salary eligibility thresholds are respectively moved to € 250,000 and € 200,000.

Figure 8: Robustness for the density functions of individual productivity for civil servants who switch to the private sector/retire vs those who stay.



Notes: As in Figure 4 except for the eligibility criteria. Eligible individuals for the 2011 policy are selected as those who received salary above €200,000 for at least one year in the period 2005-2010 and work for a public sector institution where the salary cap applies.

Figure 9: Robustness for the density functions of individual productivity for civil servants who switch to the private sector/retire vs those who stay.



Notes: As in Figure 4 except for the eligibility criteria. Eligible individuals for the 2011 policy are selected as those who received salary above € 250,000 for at least one year in the period 2005-2010 and work for a public sector institution where the salary cap applies.

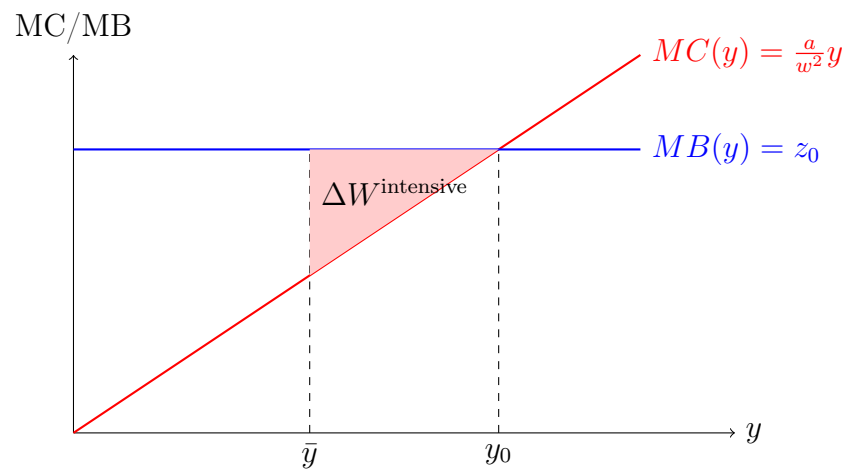


Figure 10: Deadweight-loss triangle from the salary cap.

## References

- Bellante, D. and A. N. Link (1981, April). Are public sector workers more risk averse than private sector workers? *ILR Review* 34(3), 408–412.
- Brühlhart, M., A. Eyquem, I. Z. Martínez, and E. Rubolino (2025). The output cost of inheritance. NBER Conference Paper.
- Bruni, P. (2017). The regulation of executive pay in the public and semi-public sector across the european union. Technical report, European Institute of Public Administration, Maastricht, the Netherlands.
- Buelens, M. and H. V. den Broeck (2007). An analysis of differences in work motivation between public and private sector organizations. *Public Administration Review* 67(1), 65–74.
- Buurman, M., R. Dur, and S. Van den Bossche (2012). Public sector employees: Risk averse and altruistic? *Journal of Economic Behavior & Organization* 83(3), 279–291.
- Caballero-Cordero, V., C. Ghirelli, Á. L. Gómez, and J. J. Pérez (2025). The public–private wage gap in the euro area a decade after the sovereign debt crisis. *International Journal of Manpower* 46(8), 1482–1506.
- Campos, M. M., D. Depalo, E. Papapetrou, J. J. Pérez, and R. Ramos (2017). Understanding the public sector pay gap. *IZA Journal of Labor Policy* 6(1), 7.
- Chetty, R. (2009). Is the taxable income elasticity sufficient to calculate deadweight loss? the implications of evasion and avoidance. *American Economic Journal: Economic Policy* 1(2), 31–52.
- Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica* 80(3), 969–1018.
- Chetty, R., A. Guren, D. S. Manoli, and A. Weber (2011). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review* 101(3), 471–475.
- Di Porto, E., C. Dustmann, C. Giannetto, and L. Incoronato (2025, August). Salary caps for public managers and public sector performance. Presented at the Visit-INPS Annual Conference 2025.

- Fenizia, A. (2022). Managers and productivity in the public sector. *Econometrica* 90(3), 1063–1084.
- Fenizia, A., D. Checchi, and C. Lucifora (2023). Public- and private-sector jobs: a cross-country perspective. *Oxford Economic Papers* 76(3), 759–779.
- Giordano, R., D. Depalo, M. Coutinho Pereira, B. Eugène, E. Papapetrou, J. J. Perez, L. Reiss, M. Roter, and A. Van Riet (2011). The Public Sector Pay Gap in a Selection of Euro Area Countries. ECB Working Paper 1406, European Central Bank.
- Hansen, G. D. (1985). Indivisible labor and the business cycle. *Journal of Monetary Economics* 16, 309–327.
- Huang, T. and A. R. Villadsen (2023). Top managers in public organizations: A systematic literature review and future research directions. *Public Administration Review* 83(6), 1618–1634.
- Kaplow, L. (2024). Optimal income taxation. *Journal of Economic Literature* 62(2), 637–738.
- Kleven, H., C. Kreiner, K. Larsen, and J. Søgaaard (2025). Micro versus macro labor supply elasticities: The role of dynamic returns to effort. *American Economic Review* 115(9), 2849–2890.
- Kleven, H. J. (2016). Bunching. *Annual Review of Economics* 8, 435–464.
- Kleven, H. J., C. Landais, and E. Saez (2013). Taxation and international migration of superstars: Evidence from the european football market. *American Economic Review* 103(5), 1892–1924.
- Kleven, H. J. and M. Waseem (2013). Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from pakistan. *Quarterly Journal of Economics* 128(2), 669–723.
- Lazear, E. P., K. L. Shaw, and C. T. Stanton (2015). The value of bosses. *Journal of Labor Economics* 33(4), 823–861.
- Limodio, N. (2021). Bureaucrat allocation in the public sector: Evidence from the world bank. *The Economic Journal* 131(639), 3012–3040.
- OECD (2012). *Public Sector Compensation in Times of Austerity*. Paris: OECD Publishing.

- OECD (2024). *OECD Guidelines on Corporate Governance of State-Owned Enterprises 2024*. Paris: OECD Publishing.
- Perry, J. L. and L. R. Wise (1990). The motivational bases of public service. *Public Administration Review* 50(3), 367–373.
- Piketty, T., E. Saez, and S. Stantcheva (2014). Optimal taxation of top labor incomes: A tale of three elasticities. *American Economic Journal: Economic Policy* 6(1), 230–271.
- Rainey, H. G. and P. Steinbauer (1999). Galloping elephants: Developing elements of a theory of effective government organizations. *Journal of Public Administration Research and Theory* 9(1), 1–32.
- Rogerson, R. (1988). Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics* 21(1), 3–16.
- Saez, E. (2010). Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy* 2(3), 180–212.
- Saez, E., J. Slemrod, and S. H. Giertz (2012). The elasticity of taxable income with respect to marginal tax rates: A critical review. *Journal of economic literature* 50(1), 3–50.

## Appendix

### A The choice of the dataset

As introduced in Section 3, three main datasets are potentially useful for analysing salary-cap policies in the public sector. The INPDAP dataset is not an option, as it begins only in 2014 and cannot capture pre-trends for the 2014 reform or the earlier 2011 reform. Di Porto et al. (2025) therefore use the Uniemens dataset; here, we discuss the rationale for our alternative choice of EC over Uniemens.

Uniemens contains contribution records for private-sector pension funds, primarily the *Fondo Pensioni Lavoratori Dipendenti* (FPLD). While excluding agriculture and caregivers, Uniemens was designed to eventually include all Italian employees. Over time, some public employees contributing to sectoral funds (e.g., listed companies, postal services, public transport, INPS staff) were moved into Uniemens, but the selection criteria and the timing are unclear and potentially politically driven.

In contrast, the EC randomly samples 13.2% of all Italian workers born in 1950 or later who contributed to any INPS-managed fund. Although Uniemens is larger, it is likely biased in its coverage of public managers exposed to the cap. Table 5 shows that, after excluding public companies (whose managers were exempt), 17% of public managers in EC earn above € 100,000 versus only 6% in Uniemens. This discrepancy arises because managers in ministries and public entities are not included in Uniemens. Estimates based on such a selected sample risk bias, with the direction of bias unknown, given opaque selection rules.

Comparing descriptive statistics for public administration employees (Table 2 for EC vs. Table 7 for Uniemens) confirms systematic differences: women are under-represented in Uniemens (38% vs. 60% in EC), and salaries are higher (median € 32,961 vs. € 29,444; mean € 36,162 vs. € 31,957). Because our analysis focuses on social welfare, comprehensive coverage of the public sector is essential. We therefore use EC, which ensures full representation of public administration and avoids misinterpreting public companies as representative of public employment.