



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

The Consumption and Wealth Effects of an Unanticipated Change in Lifetime Resources

Tullio Jappelli, Mario Padula

To cite this article:

Tullio Jappelli, Mario Padula (2016) The Consumption and Wealth Effects of an Unanticipated Change in Lifetime Resources. Management Science 62(5):1458-1471. <http://dx.doi.org/10.1287/mnsc.2015.2192>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2015, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

The Consumption and Wealth Effects of an Unanticipated Change in Lifetime Resources

Tullio Jappelli

University of Naples Federico II, 80138 Napoli, Italy;
Center for Studies in Economics and Finance, 80126 Napoli, Italy; and
Center for Economic and Policy Research, Washington, DC 20009, tullio.jappelli@unina.it

Mario Padula

Università della Svizzera italiana, 6900 Lugano, Switzerland;
Center for Studies in Economics and Finance, 80126 Napoli, Italy; and
Center for Economics and Policy Research, Washington, DC 20009, mario.padula@usi.ch

In 2000, Italy replaced its traditional system of severance pay for public employees with a new system. Under the old regime, severance pay was proportional to the final salary before retirement; under the new regime it is proportional to lifetime earnings. This reform entails substantial losses for future generations of public employees, in the range of €20,000–€30,000, depending on seniority. Using a difference-in-difference framework, we estimate the impact of this unanticipated change in lifetime resources, on the current consumption and wealth accumulation of employees affected by the reform. In line with theoretical simulations, we find that each euro reduction in severance pay reduces the average propensity to consume by 3 cents and increases the wealth-income ratio by 0.32. The response is stronger for younger workers and for households where both spouses are public sector employees.

Keywords: lifetime resources; consumption; wealth accumulation

History: Received August 3, 2014; accepted February 3, 2015, by Amit Seru, finance. Published online in *Articles in Advance* October 15, 2015.

1. Introduction

Recently Italy reformed the severance pay system of public employees. Before 2000, severance pay was proportional to final salary; under the new regime it is proportional to lifetime earnings. Since wages generally increase with seniority over the employment lifetime, the reform entails considerable losses for future generations of public sector retirees, in the range of €20,000–€30,000 (depending on seniority), corresponding to around one year's salary for a white-collar public sector employee. Although the context is different, the reform is similar to a switch from a social security system where benefits are proportional to the previous year's earnings, to a system where benefits are tightly linked to contributions.

This paper investigates the impact of this unanticipated change in lifetime resources on the consumption and wealth accumulation of workers affected by the reform. To address our research question, we use repeated cross-sectional data for 1989–2010, which span the prereform and postreform periods and exploit the exogenous source of variation induced by the severance pay reform. Since the reform affects only public employees, we define private employees as the “control group” and public employees as the

“treatment group,” and we compare their consumption and wealth accumulation before and after the reform. Analysis of households' responses to exogenous changes in future resources has been thoroughly studied, and has important policy implications, for instance, in relation to the impact on consumption of tax reforms (see the survey by Jappelli and Pistaferri 2010a). In this paper we identify an episode in which lifetime income changes unexpectedly, and we evaluate in a quasi-experimental setting how wealth and consumption react to such a change. The approach adopted does not require the observation of individual income shocks. Rather, we compare households that are and are not exposed to shocks, and we assume that the differences in consumption and wealth arise from realization of the shocks.

Whereas the literature on the effect of anticipated income shocks on consumption is vast, much less is known about the effect of unanticipated shocks. Few papers use a similar approach as in the present paper to identify the consumption effects of exogenous and unanticipated positive income shocks, identifying episodes where income changes unexpectedly, and evaluating in a quasi-experimental setting how consumption reacts to these changes because of unemployment or disability (Browning and Crossley 2001,

Stephens 2001) or to an unexpected income bonus (Agarwal and Qian 2014b).¹ However, the present study focuses on a shock to lifetime resources rather than a shock to current income, and it studies how consumption and wealth respond to that shock.

The paper also contributes to the literature on the consumption effect of income policies. On this front, whereas most literature focuses on the effect of expansionary policies (for instance, Parker et al. 2013), the evidence on the effect of contractionary policies is rare.² Therefore, by focusing on the consumption and wealth effect of a negative shock to lifetime resources, this paper provides complementary evidence on an important area of study, which is still underinvestigated.

In §2 we discuss several reasons why it is interesting to look at the Italian severance pay reform.³ First, severance pay represents a far larger component of households' lifetime incomes in Italy than in most other countries. Since firms annually contribute approximately 7% of their wage bill to a severance pay fund from which employees cannot withdraw until termination of their contract, this fund is currently close to 10% of gross domestic product (GDP). Second, severance pay is rather illiquid, and can be regarded as a form of forced saving by workers, who can dispose of part of their severance pay only for exceptional medical expenses or the purchase of a first dwelling. Thus, changes to severance pay legislation have an impact on the individual's earnings profile up to retirement, not just for few years. Third, changes in severance pay legislation can be regarded as an exogenous innovation in lifetime resources, providing the necessary variability to assess the impact of changes in future resources on current consumption and wealth.⁴

In §3, to guide our empirical analysis, we simulate the effect of the severance pay reform on the

wealth-income ratio and the average propensity to consume, in a life-cycle model with income uncertainty. The simulations produce trajectories of wealth and consumption income ratios before and after the reform, and show that after the reform the wealth-income ratio increases, and the consumption-income ratio falls. The simulations show also that young workers react more strongly to the reform than workers close to retirement, so we expect the former group to exhibit the largest wealth and consumption adjustments. Section 4 presents our empirical analysis. We use data from the 1989–2010 Survey of Household Income and Wealth (SHIW), a large representative survey of the Italian population carried out by the Bank of Italy. The SHIW contains detailed data on the income, consumption, wealth, and demographic characteristics of households. Using a difference-in-difference framework, our baseline estimates show that a €1 reduction in severance pay increases the wealth-income ratio by 0.32 and reduces the average propensity to consume by 0.03. Since for the average household the reform reduces lifetime income by about €23,000 (relative to the old regime), the results suggest an offset ratio (ratio of increase in wealth to reduction in lifetime income) of 0.4. We find that, as predicted by our simulations, the response is stronger for younger workers and households where both spouses are public sector employees. We also perform several robustness checks, controlling for group-specific pretreatment trends, expanding the set of control variables, and exploring the heterogeneity of the effect of the reform. Furthermore, we repeat the analysis using a matching estimator as an alternative strategy to identify the effect of the reform. Section 5 summarizes our main findings.

2. The Severance Pay Reform

Severance pay was introduced in 1927 to insure Italian employees against the risk of dismissal, but it gradually evolved into a form of deferred compensation to which the employee (public or private) is entitled, irrespective of the cause of termination of employment—whether retirement, being laid off, or quitting. For private employees, the fund is guaranteed by the national social security agency (INPS) for the case of firm failure. Severance pay has become a large component of the Italian household's lifetime income, with severance pay for workers with length of seniority in the order of three or four times their annual income.

Severance pay attracts tax benefits since it is regarded as a form of forced retirement saving. Workers can draw on part of the lump-sum payment only for exceptional medical expenses or purchase of their first home. However, this applies only to workers

¹ As discussed in Jappelli and Pistaferri (2010a), there are two other approaches to estimating the effect of income shocks on consumption. These are covariance restrictions imposed by the theory on the joint behavior of consumption and income growth as in Blundell et al. (2008), or survey question about hypothetical income changes (Shapiro and Slemrod 1995, Jappelli and Pistaferri 2014).

² An exception is Agarwal and Qian (2014b), who study a policy experiment in Singapore that reduces consumers' access to housing equity and the associated consumption effect.

³ For a description of the working of severance pay before the reform, see Brugiavini and Padula (2003).

⁴ For recent papers exploiting exogenous pension reforms to estimate the effects of changes in social security wealth on wealth or saving, see Attanasio and Brugiavini (2003), Attanasio and Rohwedder (2003), and Bottazzi et al. (2006), who generally find that shocks to future social security wealth increases current wealth, although the response is far from complete. These studies point out also that the response is heterogeneous in relation, for instance to occupation and age.

Table 1 The Severance Pay Reform

	Type of contract	Severance payment
Private employees	All contracts	Years of contributions $\times 0.0691 \times$ yearly salary. Contributions are capitalized using the $0.015 + 0.75\pi$ accrual rate.
Public employees prereform	All contracts	Years of contributions $\times 0.80 \times$ (final yearly salary/12).
Public employees postreform	Contracts signed before December 2000	Pro-rata regime, with two components. The first component is $0.8 \times$ number of years of contribution until 12/2010 \times (last yearly salary/12). The second component is $0.0691 \times$ yearly earnings, capitalized at the rate $0.015 + 0.75\pi$. The weights of the two components are given by years of service before and after December 2010.
	Contracts signed after December 2000	Years of contributions $\times 0.0691 \times$ gross yearly salary. Contributions are capitalized using the $0.015 + 0.75\pi$ accrual rate (same as for private employees).

Notes. Public sector employees are state government employees. Before the reform, a slightly different formula applied to local government employees. After the law DPCM 20/12/1999, the new regime applies to all public employees whose contract was signed after January 2001, whereas a pro-rata system applies to contract signed before 12/2001.

with more than eight years of continuous employment with the same employer. Also, withdrawals cannot exceed 70% of the severance pay accumulated at the time of the request, and at any point in time, no more than 4% of the labor force can make an early withdrawal. Severance pay benefits from a double tax advantage. First, there is a deduction from the tax base, which is determined as $P-nA$, where P is the lump-sum payment, n is the number of years of employment, and A is a constant allowance. Second, the tax rate is the average tax rate corresponding to an income of $12P/n$, which for the most senior employees is lower than their general income tax rate.⁵ Since each year employers contribute a fraction of their wage bills into a fund, from which (apart from the exceptions noted above) employees receive no pay out until termination of their employment contract, severance pay operates like an unfunded social security system.⁶

Table 1 shows that for private employees, the contribution rate to the severance pay fund is 6.91% of the gross yearly salary. Contributions are then indexed to the cost of living according to the formula $0.015 + 0.75\pi$, where π is the rate of change of the consumer price index. This implies that the return from the fund is positive for inflation rates below 6%, and negative for inflation rates above 6%, which often applied in Italy prior to the introduction of the euro. Severance pay was more generous for public than for private employees, because it was linked to the last public employment salary year, not the entire working career earnings. Since the earnings profile of public employees is relatively stable and generally increasing up to retirement based on a combination of promotions and seniority rules, the final salary almost

invariably corresponds to the highest salary received by a public employee over his or her career. The formula applied for public employees was also different: severance pay was computed as 80% of the last gross salary multiplied by the number of years of service.

The 1995 reform to the Italian social security system was aimed at reducing the imbalance between projected contributions and payouts. The reform increased the retirement age and the minimum number of contribution years for pension eligibility, and it introduced a gradual transition for both public and private employees from an earnings-based system to a contribution system.⁷ With the same aim of reducing future public payouts, in 2000 government changed the rules applying to severance pay for public employees. Under the new regime (applying to all employment contracts signed after 2000), the severance pay rules for public and private employees are the same; severance pay is linked to the working career using the same indexation formula. Table 1—severance pay reform—shows that in the transition period (contracts signed before 2000) severance pay had two components, with weights given by the number of years of contributions before and after 2000.

The switch from an earnings-based formula to a contributions-based system represents a substantial loss of wealth for public employees, especially the youngest ones or those with the shortest record of public employment. Table 2 reports the results of simple calculations based on realistic public employee earnings profiles. For a public employee retiring after 40 years of contributions, whose starting salary was €15,800 and increased at an annual real rate of 1.53%, severance pay prereform would have been €76,195 and is €58,065 after the reform. For a starting salary

⁵ While maintaining favorable tax treatment, the rules were further modified in 2001, 2008, and 2012.

⁶ National Financial Accounts show that in 2000–2010 the severance pay fund accumulated by private firms against their severance pay liabilities was in the range of 10% of GDP.

⁷ Not all workers were affected by these changes in the same way. Workers close to retirement age retained the generous prereform provisions, whereas younger workers saw their benefits substantially reduced. Attanasio and Brugiavini (2003) and Bottazzi et al. (2006) provide more detail on the pension reform.

Table 2 The Severance Pay Before and After the Reform

	After the reform		
	Before the reform	Contracts signed before December 2000	Contracts signed after December 2000
	(1)	(2)	(3)
$g = 1.53\%$, $y_0 = 15,800$	76,195	69,303	58,065
$g = 2.23\%$, $y_0 = 18,000$	116,517	100,976	77,996
$g = 2.62\%$, $y_0 = 20,000$	146,234	124,342	92,980

Notes. This table shows severance pay (in euros) for public employees before and after the reform. Severance pay is obtained assuming that employees retire after 40 years of work. The inflation rate used in the accrual rate formula is 6.5%, corresponding to the 1970–2010 historical average, g is the real yearly growth rate of earnings, and y_0 the starting salary. In the first and last rows, g and y_0 correspond to historical averages for blue- and white-collar workers. In the second row, g and y_0 correspond to the historical average for all employees. In column (2), severance pay is computed “pro-rata,” for a worker who starts working in 1995.

of €18,000 increasing at an annual real rate of 2%, the reduction in severance pay is around €38,000 (from €116,517 to €77,996 after). For a starting salary of €20,000 growing at the rate of 2.62% annually, the reduction in severance pay is of the order of €50,000 (from €146,230 to €92,980). Note that private employees were unaffected by the reform; their severance pay regime was the same before and after 2000.

The examples show that the reform reduced severance pay for public employees. The implied magnitude of these changes is substantial; for the youngest public employees (contracts signed after 2001) the reduction is between 18,000 and 54,000 euro, depending on the steepness of the earnings profile. In these examples of the severance pay reform, private employees are the control group and public employees are the treatment group, which allows us to estimate the wealth effect of the reform.

There is always a concern with natural experiments that there are unaccounted confounds, for instance, adjustments in other features of severance pay or of the wage bargain that might differentially affect the wages or other work-related benefits of public and private employees. First of all, in the private sector, severance pay contributions are simply a constant fraction of earnings (with no cap for higher earnings), and the contribution rate did not change at the time of the reform. Second, national accounts data show that between 1990 and 2010 the wage dynamics of public and private sector employees followed similar trends, although the dynamics for public employees has been more volatile, with lower growth in the 1990s and faster growth in subsequent decade. Setting nominal wages (before taxes) to 100 in 1990, in 2010 the wage of private sector employees of the industrial sector was 203.7, the wage of employees in the private service sector was 196.5, and in the public sector it was 196.2.⁸ In the microeconomic data, the wage

distribution of private and public employees does not display significant differences before and after the reforms (calculations available on request).

Another potential concern is that public sector employees might respond to the reform varying their labor supply, not just consumption or saving. Evidence from the United States on the impact of the Earned Income Tax Credit shows that small changes in some payments can cause significant labor supply effects; see Meyer (2002) and Chetty et al. (2013). In general, in the Italian public sector, overtime hours are rationed with quotas assigned to each administration to control public spending. Therefore, contractual arrangements leave small margins of adjustment along this dimension. Furthermore, in many administrations, overtime hours are considered as an extra benefit attached to the wage and are often divided equally among employees.

In households with multiple earners, spouses employed in the private sector could still respond to the severance pay reform by increasing labor supply. If household labor supply is affected, it will only make the reduction in lifetime income (and wealth) smaller. Therefore in the presence of labor supply adjustments, a positive effect of the reform on the wealth-income ratio (and a negative effect on the consumption ratio) will be more difficult to detect, but none of our conclusions would be affected.

It should also be noted that—before as well as after the reform—severance pay allows limited withdrawals in case of illness or purchase of a home. The paper takes these contractual features as given, and studies the wealth and consumption effects of a reduction in severance pay entitlements. A recent literature studies how consumers respond to the option to cash out retirement savings (Agarwal et al. 2014), the incentives to borrow from their pension plans (Lu and Mitchell 2010, Beshears et al. 2012), the effect of default options on retirement saving (Beshears et al. 2011), and whether individuals should be allowed flexible access to pension savings. We leave these

⁸ Source. ISTAT—National and Institutional Accounts Data, 1990–2010 (<http://www.istat.it>).

Downloaded from informs.org by [131.111.164.128] on 12 June 2016, at 18:41. For personal use only, all rights reserved.

implications of severance pay arrangements for future research.

3. Simulation Results

To gauge the impact of the reform on the wealth-income ratio and the average propensity to consume, we simulate a life-cycle model with isoelastic utility, finite horizon, and income uncertainty, assuming a standard income process with permanent and transitory shocks. The endogenous variables of the model are consumption and the stock of wealth (i.e., accumulated saving), as in standard life-cycle models.

In the model, we assume that severance pay is illiquid and is paid out as a lump sum at retirement age N . In the prereform regime, the severance pay of a public employee is $0.8 \times N \times Y_{N-1}$, where Y_{N-1} denotes earnings in the year before retirement; in the postreform regime, severance pay is $0.0691 \sum_{t=0}^{N-1} Y_t(1+\rho)^{N-t}$, where the accrual rate is $\rho = 0.015 + 0.75\pi$. After retirement consumers rely only on accumulated savings and severance pay to finance consumption.⁹

In the model, labor supply is fixed, and terminations of public employees are due only to retirement. With voluntary or forced terminations, severance pay is cashed before retirement. But the important feature of the reform is that regardless of when the severance pay is actually paid out, the reform reduces lifetime resources of public employees. Assuming that severance pay is paid at retirement makes the model more tractable and reflects the fact that in the public sector both voluntary and forced terminations are rare.¹⁰

To keep the model in line with the data, the simulations produce life-cycle profiles of the ratio of wealth and consumption to income, under both regimes. To simulate the model, we assume that the reform takes place unexpectedly after t^* years of work.¹¹ Comparison of the profiles in the two regimes reflects an

unanticipated negative shock to lifetime resources. Appendix A provides further details on the structure of the model and its parameterization.

The simulations indicate that an unanticipated negative income shock to lifetime resources reduces the average propensity to consume and increases the wealth-income ratio. To offset the negative shock to lifetime resources, individuals react by reducing their consumption and increasing their stock of wealth through increased savings. The response of consumption and wealth is a standard implication of life-cycle models when some components of future resources, for instance social security wealth, are reduced. The crucial assumption of the model is that individuals smooth consumption over time. When the negative shock to severance pay is announced, consumption falls and saving increases to allow individuals to build up enough wealth to prevent consumption from falling too much when the reduction of severance pay actually realizes. The effects on consumption and wealth depend on the size of the shock, and therefore are stronger for younger workers; for them, the change in the severance pay regime applies for a longer time span, and therefore the reduction in future resources is larger.

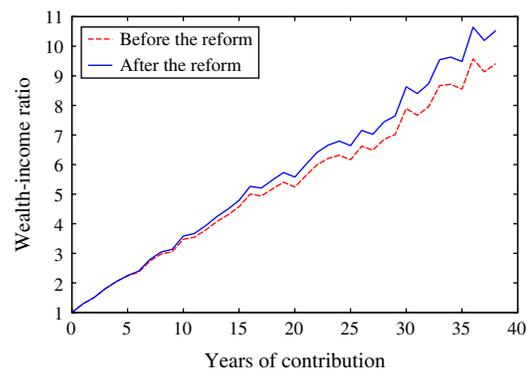
To provide an indication of the size of these potential effects, Figure 1 plots the simulated profiles of the wealth-income ratio against years of work. The lower line represents the wealth-income ratio under the old regime; the upper line is the wealth-income ratio of a consumer who experienced the reform in her fifth working year ($t^* = 5$). It is apparent that the wealth effect of the reform gradually builds over the consumer's lifetime. Figure 2 provides further insights into the effect of the reform by comparing wealth trajectories for different values of t^* (5, 15, 25, and 35). Figure 2 plots the difference between the postreform wealth profile and the baseline profile, that is, the

⁹ The simulated consumption and wealth effects are quite similar if social security contributions are proportional to earnings and benefits proportional to lifetime income.

¹⁰ We consulted available statistics on forced and voluntary terminations. Data from the Ministry of Economics and Finance (2014) indicate that in 2012–2013 forced terminations (layoffs) represented only 0.016% of the total number of public employees, including state and local public administrations (about 500 cases out of 3.3 million employees). Data for earlier years are not available. Voluntary terminations are more difficult to identify from official statistics. From the same source, we guess that they represent between 15,000 and 32,000 cases each year, or between 0.5% and 1% of the total number of public employees.

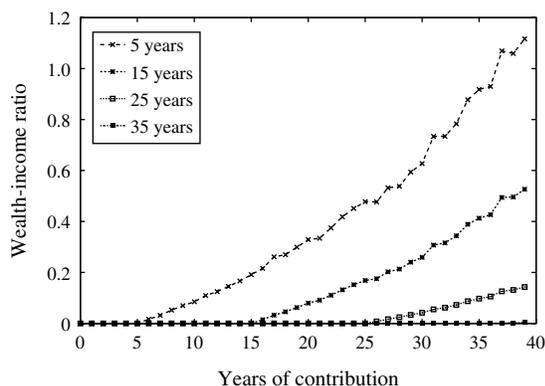
¹¹ We assume that the growth rate of real earnings equals the 1970–2010 average for the Italian economy (2.3%), the real interest rate is 1.5%, and the coefficient of relative risk aversion is 2. The standard deviations of permanent and transitory shocks are 0.16 and 0.28, respectively, as in Jappelli et al. (2008). The inflation rate used in the accrual rate formula is 6.5% (the average inflation rate in the 10 years before the reform), and retirement is set at 40 working years.

Figure 1 (Color online) The Simulated Wealth-Income Ratio Before and After the Severance Pay Reform



Notes. This figure plots the simulated wealth-income ratio before the severance pay reform (lower curve) and after the reform (upper curve). Appendix A reports the parameters used in the simulations.

Figure 2 Simulated Change in the Wealth-Income Ratio After the Severance Pay Reform

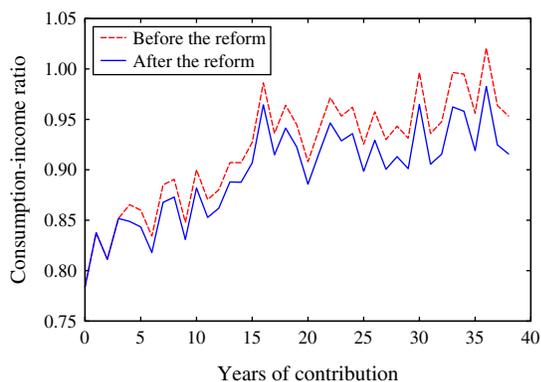


Notes. This figure plots the change in the simulated wealth-income ratio before and after the severance pay reform for four different groups of workers. The first group experiments the reform after five years of work (top curve), the second group after 15 years (second from the top), the third after 25 years (third from top), and the fourth group after 35 years (bottom curve). Appendix A reports the parameters used in the simulations.

wealth profile that would be observed in the absence of reform. The wealth effect is positive in all cases but stronger the earlier it occurs in the employee's career. Thus, a worker experiencing reform at $t^* = 5$ faces a much bigger reduction in lifetime resources than someone close to retirement. Figures 3 and 4 show that the reform reduces the average propensity to consume. The reduction is in the 1%–3% range, and is larger for workers who experience the shock earlier in their working life ($t^* = 5$ and $t^* = 15$ in Figure 4) relative to those close to retirement ($t^* = 35$).

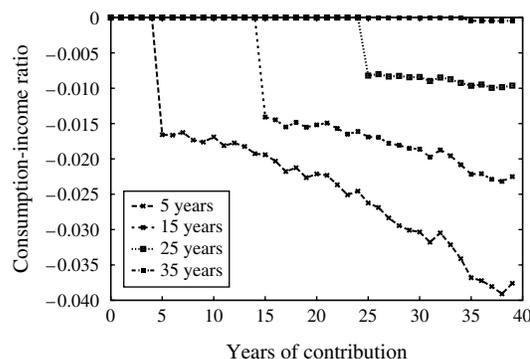
Overall, the simulations suggest that an unanticipated negative income shock to lifetime resources reduces the average propensity to consume and increases the wealth-income ratio. Furthermore, both effects depend on the size of the shock, and there-

Figure 3 (Color online) The Simulated Consumption-Income Ratio Before and After the Severance Pay Reform



Notes. This figure reports the simulated consumption-income ratio before the severance pay reform (upper curve) and after the reform (lower curve). Appendix A reports the parameters used in the simulations.

Figure 4 Simulated Change in the Consumption-Income Ratio After the Severance Pay Reform



Notes. This table reports the change in the simulated consumption-income ratio before and after the severance pay reform of four different groups of workers. The first group experiments the reform after five years of work (bottom curve), the second group after 15 years (second from the bottom), the third after 25 years (third from bottom), and the fourth group after 35 years (top curve). Appendix A reports the parameters used in the simulations.

fore are stronger for younger workers. To test these theoretical predictions, we now turn to the empirical analysis.

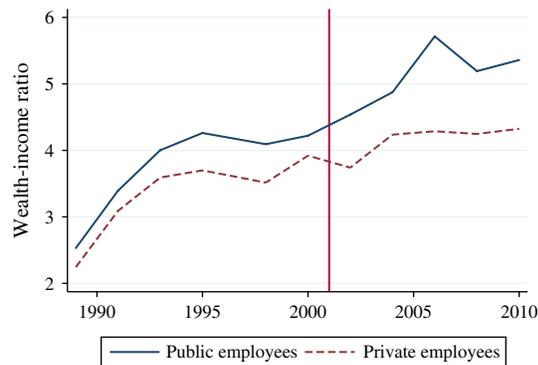
4. Empirical Estimates

Table 3 reports sample statistics for public and private employees in the pooled 1989–2010 sample, a total of 28,665 observations. We restrict the sample to households where the household head is aged 20–55, and is employed in either the public or the private sector, thereby excluding self-employed people (who of course are not entitled to severance pay) and workers near to retirement. The sample includes 61% private employees and 39% public employees. However, 48% of households have at least one public sector employee—the household head, the spouse, or both. We exploit this information to check whether the

Table 3 Sample Statistics for Variable Used in the Estimation

Variable	Mean	Median	SD
Public employee	0.39	0.00	0.49
Public employee in the household	0.48	0.00	0.50
Private employee	0.61	1.00	0.49
Age	42.16	43.00	7.95
Male	0.79	1.00	0.41
College degree	0.12	0.00	0.33
High school	0.39	0.00	0.49
Family size	3.25	3.00	1.28
Resident in the north	0.47	0.00	0.50
Resident in the centre	0.21	0.00	0.41
Resident in the south	0.32	0.00	0.47
Wealth/income	3.88	3.19	4.19
Consumption/income	0.77	0.74	0.59
Disposable income	28.69	25.65	16.38

Notes. Data are drawn from the 1989–2010 SHIW. The sample includes 28,665 observations.

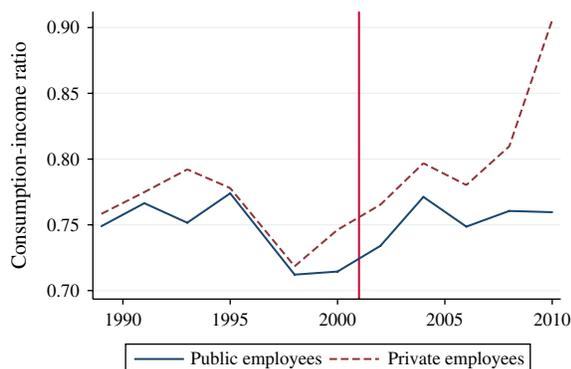
Figure 5 (Color online) Wealth-Income Ratio, by Occupation

Notes. This figure shows the time-series profile of the median wealth income ratio by occupation group. The continuous line refers to public employees, the dashed to private employees. Data are drawn from the 1989–2010 SHIW.

reform has a stronger effect for households with more than one public sector employee, and to select a sample of households with at least one public employee.

Net wealth is the sum of net financial assets and real assets. Net financial assets is the sum of transaction accounts, government bonds, CDs, corporate bonds, retirement accounts, life insurance, and stocks, less household debt. Real assets are the sum of real estate, unincorporated business holdings, valuables, and art objects. Consumption is measured as non-durable expenditures. Appendix B provides further information about the survey.

For the whole sample, the wealth-income ratio is 3.88 (median is 3.19) and the consumption-income ratio is 0.77 (median is 0.74). The ratios differ by employment group and exhibit different trends. Figures 5 and 6 show that private employees have a lower wealth-income ratio and a higher consumption-income ratio than public employees, both before and after the reform. Notice that in Figures 5 and 6 the gap between the two lines widens after the reform.

Figure 6 (Color online) Consumption-Income Ratio, by Occupation

Notes. This figure shows the time-series profile of the median wealth income ratio by occupation group. The continuous line refers to public employees, the dashed to private employees. Data are drawn from the 1989–2010 SHIW.

Table 4 Wealth-Income and Consumption-Income Ratios Before and After the Severance Pay Reform

	Prereform	Postreform	Change after the reform
Wealth-income ratio			
Private employees	3.29	4.17	0.88
Public employees	3.71	5.11	1.40
Difference	0.42	0.94	0.52
Consumption-income ratio			
Private employees	0.76	0.81	0.05
Public employees	0.75	0.75	0.00
Difference	−0.01	−0.06	−0.05

Notes. Data are drawn from the 1989–2010 SHIW. The sample includes 28,665 observations.

Table 4 shows that the difference in the wealth-income ratio increases from 0.42 before the reform to 0.94 after the reform. The difference-in-difference estimates are 0.52 for the wealth-income ratio and −0.05 for the consumption-income ratio, and both are statistically different from zero at the 1% level. These estimates show that after the reform public employees have increased their wealth and reduced their consumption relative to private employees, which confirms the simulation analysis.

For several reasons, the evidence provided in Table 4 is not conclusive about the effect of the reform. First, it does not consider that other variables (such as age, education, income) might shift wealth and consumption ratios after the reform. Macro shocks may also affect the two variables differently over time; examples from the early 2000s include the stock market crash and subsequent recovery, the decline in yields from short-term government bonds after the introduction of the euro, and the house price boom.

To address these issues, we rely on regression analysis, which allows us to study the exogenous variation in lifetime income brought about by the reform, controlling for households' characteristics and group-specific trends. We test for the effect of the reform using the following regression framework:

$$y_{it} = \alpha + \beta M_i + \gamma POST_t + \delta M_i \times POST_t + \theta x_{it} + \varepsilon_{it},$$

where y is the ratio of wealth or consumption to disposable income, M is a dummy for the treatment group (public employees), $POST$ is a dummy for the postreform period (2002–2010), x a vector of the control variables (age, gender, education, family size), and ε is an error term.¹² The parameter δ measures the effect of the reform. According to our simulation

¹² Note that by appropriate redefinition of the variable M and the treatment group, this framework could be extended to examine the differential impact of the reform on households with more than one public employee (i.e., both spouses work in the public sector) or specific population groups.

analysis, we expect $\delta > 0$ in the regression for the wealth-income ratio, and $\delta < 0$ in the regression for the consumption-income ratio.

The validity of our estimates rests on two assumptions: (1) the severance pay reform is exogenous with respect to consumer decisions, and (2) the reform is exogenous with respect to changes in sample composition. In relation to assumption (1), the possible endogeneity of the reform can be ruled out. The 2000 reform was not implemented in order to offset the different wealth accumulation paths of the employment groups; rather, it was part of a deficit-reduction package aimed at reducing projected outlays in the public sector.

Assumption (2) requires that the reform does not cause changes in the sample composition. These are possible if labor supply of public employees changes after the reform. The dynamics of labor income of private and public employees shows that the difference between the two groups is small, and that trends are similar over the sample period (see the online appendix, available as supplemental material at <http://dx.doi.org/10.1287/mnsc.2015.2192>). As noted in §2, even if the severance reform does have any significant impact on labor supply, it will make the reduction in lifetime income and wealth smaller, which would make our results even stronger.¹³ In addition, to assess the validity of assumption (2), we check if job mobility from public to private employment (and vice versa) is independent of the severance pay reform, that is, that workers did not switch jobs as a result of the reform. Since the SHIW has a rotating panel component, we can check the validity of this assumption by computing transition rates across the two employment groups for each pair of adjacent survey years from 1989 to 2010. We find that, in each period, the probability of not changing sector is about 90% for both groups. Furthermore, we do not reject the hypothesis that the degree of sector mobility is the same before and after the reform for each of the estimated transition matrices, even controlling for household characteristics. Although we cannot directly test the hypothesis that workers did not change sectors as a consequence of the reform, we take this as indirect evidence that the severance pay reform has not changed the overall pattern of worker mobility.

4.1. Baseline Estimates

Table 5 reports the baseline estimates for the effect of the 2000 reform on the wealth-income ratio. The

¹³ Note that income changes unrelated to the reform cannot simultaneously explain the positive effect on the wealth-income ratio and the negative effect on the consumption-income ratio. For instance, if the reform has no effect on wealth and consumption and income increases after the reform for public more than for private employees, both the wealth and the consumption-income ratios would fall.

Table 5 Baseline Specifications

	Wealth-income ratio		Consumption-income ratio	
	(1)	(2)	(3)	(4)
<i>Public employee</i>	0.417 (0.065)***	0.015 (0.065)	-0.017 (0.009)*	0.002 (0.010)
<i>Postreform period</i>	0.881 (0.063)***	0.772 (0.063)***	0.048 (0.009)***	0.059 (0.009)***
<i>Public employee × Postreform</i>	0.526 (0.103)***	0.321 (0.101)***	-0.040 (0.015)***	-0.030 (0.015)**
<i>Age</i>		0.088 (0.003)***		-0.004 (0.000)***
<i>Male</i>		0.093 (0.063)		-0.009 (0.009)
<i>Family size</i>		0.051 (0.021)**		-0.009 (0.003)***
<i>College degree</i>		1.814 (0.079)***		-0.137 (0.012)***
<i>High-school diploma</i>		1.290 (0.052)***		-0.082 (0.008)***
<i>Resident in the centre</i>		0.507 (0.063)***		0.037 (0.009)***
<i>Resident in the south</i>		-0.040 (0.057)		0.092 (0.008)***
<i>Constant</i>	3.291 (0.041)***	-1.231 (0.147)***	0.763 (0.006)***	0.978 (0.021)***
<i>R²</i>	0.02	0.08	0.00	0.02

Notes. Data are drawn from the 1989–2010 SHIW. Standard errors are reported in parentheses. The sample includes 28,665 observations.

***Denotes statistical significant at 1% confidence level; **significant at 5% level; *significant at 10% level.

positive coefficient of the public employment dummy mirrors the difference between employment groups in Figure 5 and shows that the wealth-income ratio is 0.42 higher for public employees than private employees. The positive coefficient of the postreform dummy indicates the existence of a common trend since the wealth-income ratio of both groups increases by 0.88. The positive coefficient of the interaction term between the postreform dummy and the treatment group indicates that the reform has increased the wealth-income ratio of public employees relative to private employees by 0.53.

The second regression in Table 5 also includes demographics controls.¹⁴ Age and education are proxies for lifetime earnings, and regional dummies control for differences in wealth across Italian macroregions. The coefficients of these additional variables have the expected sign. Wealth increases with age (equivalent to about one year's earnings every 10 years) and is higher for households headed by males who are high-school or college graduates. The

¹⁴ The reference group is a private employee, without a college or high-school degree, living in northern Italy.

parameter δ is 0.32 and is quite precisely estimated, showing that the impact on wealth of the severance pay reform is about one-third of the disposable income; evaluated on average disposable income, this corresponds to an impact of €9,180 or four months salary. Since the calculations in Table 2 show that for a public employee who enters the labor market after 2000 (expecting a growth rate of earnings of 2.2%) the reform has reduced lifetime income by €22,980 (relative to the old regime), the result suggests an offset ratio—the ratio between the increase in wealth and the reduction in severance pay—of 0.4. In a different context, this value is not far from estimates of the offset rate between social security wealth and private wealth; see, for instance, Gale (1998), Attanasio and Brugiavini (2003), and Bottazzi et al. (2006).

The other two regressions in Table 5 focus on the consumption-income ratio.¹⁵ They show that public employees average propensity to consume reduced after the reform by four percentage points (column (3)) and three percentage points (column (4)). Note also that the positive coefficient of the postreform dummy in the consumption regression reflects an aggregate effect, that is, a generalized, across-the-board increase in the consumption-income ratio for all groups, both those affected and those unaffected by the reform. The increase in the propensity to consume in the regressions of Table 5 matches a general Italian trend of the past two decades, with a substantial decline of the average propensity to consume of the household sector (over 10 percentage points from 1990 to 2010). The regressions of Table 5 show that for public employees the increase in the consumption-income ratio after the reform (the interaction term) is not as strong as for private employees, which we attribute to the reform. The results of Table 5 rest on a number of assumptions, which are discussed in §4.3, where we perform several robustness checks and investigate the presence of group-specific pretreatment trends, the robustness to alternative definitions of the treatment group, and the heterogeneity of the effect of the reform.

In Table 6 we redefine the control group, considering as “treated” all households with at least one member (not necessarily the household head) who is a public sector employee. The effects are similar to those presented in Table 5. In particular, the reform increases the wealth-income ratio by 0.3 (column (2)), and reduces the consumption-income ratio by 0.024 (column (4)).

Since we are comparing two decades of data (pre and post 2000), the jump in the ratio is not what we

¹⁵ The response of durable consumption might be different than that of nondurables. In the simulation, consumption refers to nondurable expenditures, and therefore we prefer to use this definition of consumption also in the regressions.

Table 6 At Least One Public Employee in the Household

	Wealth-income ratio		Consumption-income ratio	
	(1)	(2)	(3)	(4)
<i>Public employee in the household</i>	0.421 (0.063)***	−0.048 (0.063)	−0.057 (0.009)***	−0.033 (0.009)***
<i>Postreform period</i>	0.863 (0.069)***	0.742 (0.068)***	0.045 (0.010)***	0.056 (0.010)***
<i>Public employee in the household × Postreform</i>	0.466 (0.100)***	0.306 (0.098)***	−0.030 (0.014)**	−0.024 (0.014)*
<i>Age</i>		0.088 (0.003)***		−0.004 (0.000)***
<i>Male</i>		0.069 (0.063)		−0.012 (0.009)
<i>Family size</i>		0.049 (0.021)**		−0.008 (0.003)***
<i>College degree</i>		1.843 (0.079)***		−0.125 (0.012)***
<i>High-school diploma</i>		1.298 (0.053)***		−0.076 (0.008)***
<i>Resident in the centre</i>		0.513 (0.063)***		0.039 (0.009)***
<i>Resident in the south</i>		−0.024 (0.057)		0.093 (0.008)***
<i>Constant</i>	3.249 (0.045)***	−1.205 (0.147)***	0.785 (0.006)***	0.981 (0.022)***
<i>R</i> ²	0.02	0.08	0.00	0.02

Notes. Data are drawn from the 1989–2010 SHIW. Standard errors are reported in parentheses. The sample includes 28,665 observations.

***Denotes statistical significant at 1% confidence level; **significant at 5% level; *significant at 10% level.

observe in the year after the reform, but rather the increase of the conditional mean of the ratio in 2002–2010 compared to 1989–2000. To distinguish between the short- and long-run responses of wealth and consumption after the reform, we run regressions on two subsamples, one including only the years 2000 and 2002 (just before and after the reform), and the other including only the first and the last years of data (1989 and 2010). The results, reported in the online appendix, show a stronger effect of the reform in 1989 and 2010 than in the 2000–2002 sample. Note, however, that the sample size is smaller than in our baseline regressions, so standard errors tend to be larger.

4.2. Tests by Number of Public Employees and Years of Contributions

In Table 7 we distinguish between households with only one public employee (one of the two partners or another family member) and households with two or more public employees. Since the wealth loss induced by the reform is larger if there is more than one public employee in the household, we expect the reform to have a stronger effect on these households. In the wealth-income ratio regression, the coefficient of the

Table 7 Distinguishing Between One and More Than One Public Employee in the Household

	Wealth-income ratio		Consumption-income ratio	
	(1)	(2)	(3)	(4)
<i>One public employee</i>	0.361 (0.069)***	−0.001 (0.068)	−0.032 (0.010)***	−0.016 (0.010)*
<i>More than one public employee</i>	0.598 (0.100)***	−0.204 (0.101)**	−0.130 (0.014)***	−0.090 (0.015)***
<i>Postreform period</i>	0.863 (0.069)***	0.745 (0.068)***	0.045 (0.010)***	0.058 (0.010)***
<i>One public employee × Postreform</i>	0.454 (0.108)***	0.281 (0.105)***	−0.032 (0.015)**	−0.026 (0.015)*
<i>More than one public employee × Postreform</i>	0.537 (0.166)***	0.363 (0.161)**	−0.034 (0.024)	−0.030 (0.024)
<i>Age</i>		0.088 (0.003)***		−0.004 (0.000)***
<i>Male</i>		0.078 (0.063)		−0.008 (0.009)
<i>Family size</i>		0.054 (0.021)**		−0.006 (0.003)*
<i>College degree</i>		1.871 (0.080)***		−0.113 (0.012)***
<i>High-school diploma</i>		1.306 (0.053)***		−0.072 (0.008)***
<i>Resident in the centre</i>		0.513 (0.063)***		0.039 (0.009)***
<i>Resident in the south</i>		−0.026 (0.057)		0.093 (0.008)***
<i>Constant</i>	3.249 (0.045)***	−1.240 (0.148)***	0.785 (0.006)***	0.965 (0.022)***
<i>R²</i>	0.02	0.08	0.01	0.02

Notes. Data are drawn from the 1989–2010 SHIW. Standard errors are reported in parentheses. The sample includes 28,665 observations.

***Denotes statistical significant at 1% confidence level; **significant at 5% level; *significant at 10% level.

interaction term between the postreform dummy and the dummy “one public employee” is 0.28, and the interaction coefficient of the dummy for “more than one public employee” is 0.36. In the consumption regressions, the variable “more than one public employee × post-reform” shows that households with more than one public employee reduced consumption by more than those with only one public employee (−0.34 versus −0.32 in column (3) and −0.30 versus −0.26 in column (4)). However, the coefficients of these regressions are estimated with larger standard errors than in the baseline regressions.¹⁶

The simulation analysis in §3 suggests that the reform has the strongest impact on young public employees since the reduction in severance pay is relatively small for employees close to retirement. To test

¹⁶ In the online appendix, we repeat the estimation distinguishing between a treated group in which all family members are public sector employees and a control sample that has all family members as private sector employees. Results are qualitatively similar to the baseline regressions. See §4.3.

this prediction, Table 8 reports the regressions for the wealth-income ratio splitting the sample by number of years of contributions of the household head. We find that the coefficient of the interaction term is 0.41 for workers with less than 30 years of contributions and is not statistically different from zero for workers close to retirement (more than 30 years of contributions).

The results for the consumption-income ratio, reported in columns (3) and (4) of Table 8, show that the interaction coefficient is larger in absolute value for younger workers (−0.034) than for older workers, for whom the coefficient is not precisely estimated...¹⁷

4.3. Robustness Checks

The main challenge to the identification design in a difference-in-difference framework is the potential effect of group-specific pretreatment trends. If

¹⁷ In the online appendix we report estimates using age splits, showing that the effect is mostly concentrated among the young (less than 45 years old). See §4.3.

Table 8 Sample Splits by Years of Contributions to the Severance Pay Fund

	Wealth-income ratio		Consumption-income ratio	
	≤ 30 (1)	> 30 (2)	≤ 30 (3)	> 30 (4)
<i>Public employee</i>	−0.078 (0.063)	0.392 (0.150)***	0.006 (0.007)	−0.022 (0.010)**
<i>Postreform period</i>	0.672 (0.071)***	1.194 (0.152)***	0.067 (0.016)***	0.018 (0.010)*
<i>Public employee × Postreform</i>	0.406 (0.119)***	−0.081 (0.253)	−0.034 (0.014)**	−0.011 (0.016)
<i>Age</i>	0.097 (0.004)***	0.030 (0.027)	−0.004 (0.001)***	−0.010 (0.003)***
<i>Male</i>	0.059 (0.073)	0.430 (0.180)**	−0.005 (0.011)	−0.020 (0.018)
<i>Family size</i>	0.071 (0.024)***	−0.128 (0.053)**	−0.008 (0.005)*	−0.019 (0.004)***
<i>College degree</i>	1.760 (0.086)***	1.994 (0.358)***	−0.151 (0.010)***	−0.060 (0.016)***
<i>High-school diploma</i>	1.263 (0.057)***	1.293 (0.140)***	−0.094 (0.010)***	−0.028 (0.009)***
<i>Resident in the centre</i>	0.561 (0.072)***	0.287 (0.141)**	0.036 (0.020)*	0.045 (0.008)***
<i>Resident in the south</i>	−0.027 (0.064)	−0.151 (0.141)	0.086 (0.008)***	0.119 (0.011)***
<i>Constant</i>	−1.558 (0.160)***	1.741 (1.389)	0.954 (0.040)***	1.297 (0.153)***
Observations	23,767	4,898	23,767	4,898
R ²	0.08	0.06	0.01	0.04

Notes. This table reports ordinary least squares regressions for sample splits defined by the number of years of contributions. Data are drawn from the 1989–2010 SHIW. Standard errors are reported in parentheses.

***Denotes statistical significant at 1% confidence level; **significant at 5% level; *significant at 10% level.

the wealth-income ratio and the consumption-income ratio evolve differently between public and private employees before the reform, our results would falsely detect an effect of the reform when the effect can actually not be there. To control for the existence of pretreatment trends, we follow Bell et al. (1999) and perform two checks. The first amounts to restrict our sample to the years before the reform and to redefine the postreform dummy as a variable taking value 1 after 1995 and value 0 otherwise, which means to pretend that the reform has taken place in 1995.¹⁸ In the second check we retain the whole sample (years 1989–2010), but add to the baseline specification the post-1995 dummy and its interaction with the public employee dummy.¹⁹

¹⁸ Notice that there are two surveys before the reform, 1995 and 1998. Thus, using 1995 as “reform year” allows using two years of data before the reform (1989 and 1991) and two years after (1998 and 2000).

¹⁹ Group-specific trends are problematic for the differences-in-differences design both before and after the reform. We also check the stability of the results removing post-2008 observations. The results are similar to those reported in Table 5.

In both checks, the main coefficient of interest is that of the interaction between the public employee and the post-1995 dummies, which would be statistically different from zero in the presence of group-specific prereform trends. The results are reported in the online appendix, and show that the coefficients of the interaction term are not statistically significant, thus providing no support for the existence of group-specific pretreatment trends (columns (1)–(4)). Moreover, columns (3) and (6) show that the coefficients of the interaction term between the public employee and the “true” postreform dummies are statistically significant and comparable in size to those found in Table 5, after one controls for pretreatment group-specific trends through the interaction between the public employee and the post-1995 dummies.

A second challenge to our identification design is related to how the treatment group is defined. We consider the treated group households whose head is a public employee. Therefore, it may well be that the treatment group contains households whose head is a public employee while other members are not. In the online appendix, we redefine the treatment group

as households where all members are public employees and the control group as households where all members are private employees (other households are excluded from the analysis). The results, reported in columns (1) and (2) for the wealth-income ratio and in columns (3) and (4) for the consumption-income ratio, are similar to those reported in Table 5, thus supporting the validity of our baseline definition of treatment and control groups.

The identification assumption of difference-in-difference estimates are more credible through the conditioning on a set of control variables, and for this reason, in Table 5, we include age, gender, family size, education, and area of residence. Adding these controls attenuates the effects of the reform but does not alter our main conclusions. However, one still wonders if one can make the identification assumptions even more credible by expanding the set of controls. This is achieved in the online appendix, where we opt for a finer description of the effects of geography, obtained by replacing the area dummies with a finer classification based on 19 regional dummies. The results suggest that using regional dummies does not affect the main conclusions on the effect of the reform on the wealth-income and consumption-income ratios. In the online appendix, we add sector dummies to the baseline specification, and results are also similar to the baseline estimates.

Another threat to the differences-in-differences design has to do with the heterogeneity of the treatment effects. The baseline results in Table 5 rely on the standard difference-in-difference hypothesis by assuming that the effect of the reform is homogeneous. However, if the true effect is heterogeneous along some relevant dimension, the estimated effect is some average of the underlying effects, and except for special cases, might bear little resemblance with them. The simulations show that the effects are indeed heterogeneous along a relevant dimension, and in particular that they are with the effect being larger for younger households. Table 8 already investigates this implication, which is further addressed in the online appendix, which reports sample splits by age. The results indicate a stronger effect for households headed by individuals younger than 45 years, in line with our expectations, and imply that the results for the whole sample are mainly driven by the group of younger households.

Although not being explicitly rooted in our simulation model, education is a further dimension of heterogeneity, which can be relevant for our results. The online appendix documents the heterogeneity of the effect between education groups. The wealth response of the group with higher education is stronger than that of the group with less education, a fact that

might be linked to financial sophistication of people with higher education. However, the consumption response is similar in the two groups. Therefore, we abstain from making strong conclusions from the group estimates.

A final concern is that the difference-in-difference estimator requires that the expectations of the potential outcomes (the wealth-income ratio and the consumption-income ratio) do not depend on the treatment status (public versus nonpublic employee), conditional on a set of covariates; that is, differences-in-differences estimators are not valid in the presence of pretreatment trends. Matching estimators provide an alternative identification strategy. However, since they do not allow controlling for selection on unobservable variables, one can view matching as an alternative approach to identify the effect of the reform. We therefore also rely on results obtained using the propensity score matching and the nearest neighbor matching algorithm. The propensity score is estimated using age, gender, family size, education, and geographical area as controls and satisfies the balancing property at standard significance levels. Results are reported in the online appendix and are quite similar to those reported in the main text.

5. Summary

We study how an unanticipated negative shock on lifetime resources affects households' consumption and wealth. The negative income shock we consider is the 2000 Italian severance pay reform, which has resulted in a significant wealth loss for public employees but does not affect private employees. Therefore, the reform provides the quasi-experimental setting to identify the effect of a negative income shock on consumption and wealth.

To gauge the impact of the shock, we simulate a standard life-cycle model of intertemporal choice with income uncertainty and a parameterization of severance pay that closely resembles the Italian prereform and postreform regimes. The simulations show that the shock reduces the average propensity to consume and increases the wealth-income ratio. Furthermore, since the shock is greater for young individuals, both responses are larger for individuals with longer retirement horizons relative to those close to retiring.

Relying on data from the Bank of Italy SHIW, we used a large representative sample of the Italian population available for 1989 to 2010 to show Italian households' responses to the reduction in future severance pay brought about by the reform—accumulation of more private wealth and reduced consumption. In our baseline estimates, a reduction in severance pay equal to one year's income is followed by an increase in wealth of about four months

income and a reduction of three percentage points in the average propensity to consume. The empirical analysis yields two other results that are in line with our simulation analysis: (i) the wealth response is stronger among households with more than one public employee, and (ii) the effect of the reform is stronger for young workers, who expect the strongest decline in severance pay.

Supplemental Material

Supplemental material to this paper is available at <http://dx.doi.org/10.1287/mnsc.2015.2192>.

Acknowledgments

The authors thank Annalisa Scognamiglio and seminar participants at the Conference on Policy Implications of Demographic Aging (Padua, November 18–19, 2014) for comments. The authors also thank Elena Crivellaro for excellent research assistance. T. Jappelli acknowledges financial support from the European Commission's Seventh Framework Programme [Project MOPACT, Grant FP7-SSH-2012-1/320333]. M. Padula acknowledges financial support from the Italian Ministry of Universities and Research [FIRB project on the Economic Effects of Demographic Aging].

Appendix A. Simulation of the Consumption and Wealth Effects of the Reform

To simulate the effect of the severance pay reform on the propensity to consume and the wealth-income ratio, we assume that consumers have finite horizons and choose consumption, C_t , to maximize the following objective function:

$$\max \left\{ E_0 \sum_{t=0}^{T-1} \beta^t U(C_t) \right\},$$

where β is the subjective discount factor, the instantaneous utility function is $(C_t^{1-\gamma} - 1)(1 - \gamma)^{-1}$, and $\gamma > 0$ is the coefficient of relative risk aversion. The dynamic budget constraints are

$$A_{t+1} = (1+r)(A_t + Y_t - C_t), \quad \text{for } \tau = 0, \dots, N-1,$$

$$A_{t+1} = (1+r)(A_t + Y_t + S - C_t), \quad \text{for } \tau = N,$$

$$A_{t+1} = (1+r)(A_t + Y_t - C_t), \quad \text{for } \tau = N+1, \dots, T-1,$$

where r , Y , S , N are, respectively, the interest rate, net earnings, severance pay, and retirement age. We assume that until retirement a public employee faces the following income process:

$$Y_{t+1} = P_{t+1} V_{t+1},$$

$$P_{t+1} = GP_t Z_{t+1},$$

where G is the growth rate of earnings, P_{t+1} is the permanent component of earnings, and V_{t+1} and Z_{t+1} are i.i.d. shocks with mean equal to 1. Severance pay is illiquid and is paid out as a lump sum at age N . In the prereform regime, the severance pay of a public employee is

$$0.8 \times N \times Y_{N-1}.$$

In the postreform regime, severance pay is computed as

$$0.0691 \sum_{t=0}^{N-1} Y_t (1+\rho)^{N-t},$$

where the accrual rate is $\rho = 0.015 + 0.75\pi$. After retirement, consumers rely only on accumulated savings and severance pay to finance consumption. Notice, however, that the simulated consumption and wealth effects are quite similar if one introduces social security contributions proportional to earnings and benefits proportional to lifetime income.

We solve the model using the endogenous grid point algorithm and exploiting the homogeneity of the utility function to express the variables as a ratio of the permanent component of income. The simulations produce life-cycle profiles of consumption and cash on hand (wealth plus income) in both regimes. Note that the simulated profiles in the old regime describe the consumption and cash-on-hand trajectories had the reform not taken place.

To simulate the model, we assume that the reform takes place unexpectedly after t^* years of work. We assume that the growth rate of real earnings equals the 1970–2010 average growth rate of earnings for the Italian economy (2.3%), the real interest rate is 1.5%, and the coefficient of relative risk aversion is 2. The standard deviations of permanent and transitory shocks are 0.16 and 0.28, respectively, as in Jappelli et al. (2008). The inflation rate used in the accrual rate formula is the average inflation rate in 1991–2000 (6.5%) and retirement age is set at 40 working years. We simulate the model for 1,000 individuals and report average consumption and wealth profiles in Figures 1 through 4.

Appendix B. The SHIW

In Italy, the main source of microeconomic data for macroeconomists is the Bank of Italy SHIW, which collects detailed information on demographics, households' consumption, income, real wealth, and financial wealth. The availability of household data on income, consumption, and wealth in the same data set (as well as the presence of a panel component) is what makes the SHIW a unique reference for researchers interested in consumption, income, and wealth inequality, and their changes over time. Jappelli and Pistaferri (2010b) present stylized facts on labor supply, income, consumption, wealth, and several measures of consumption and income inequality.

The SHIW was conducted annually from 1965 to 1987, and then biannually since 1987 (with the exception of a three-year interval between 1995 and 1998). In 1987, the SHIW oversampled rich households to collect more accurate statistics on asset holdings. Consumption and income data are available for all years, but information on wealth is available only since 1987 (before 1987, only information on real assets is available). In the paper, we use data for 1989–2010.

The SHIW is a representative sample of the Italian resident population. The sample design is consistent with that used by the Labor Force Survey conducted by ISTAT, and therefore the participation rate, and the distribution of employees by sector and occupation match those of the general population. Sampling is carried out in two stages: the first consists of the selection of municipalities, the second the selection of households. Municipalities are categorized into 51 strata, defined by 17 regions and three classes

of population size (over 40,000, 20,000–40,000, less than 20,000). All municipalities in the first group are included. Those in the second and third groups are selected randomly with a probability proportional to their population size. In the second stage, households are selected randomly from registry office records. The response rate has been between 35 and 50% throughout the sample period.

Data are collected through personal interviews in the first months of the calendar year (February to May). Flow income and consumption refer to the previous year, which in Italy coincides with the calendar year. Wealth and debt variables are end-of-period values. Questions concerning the whole household are answered by the head of the family or by the person most knowledgeable with the family finances. Questions on individual incomes are answered by each member, wherever possible. The unit of observation is the family, which is defined to include all persons residing in the same dwelling who are related by blood, marriage, or adoption. Individuals selected as “partners or other common-law relationships” are also treated as families.

Details about the survey are published in the Bank of Italy official website. Publications reporting the main results of each survey and the questionnaires are available at http://www.bancaditalia.it/statistiche/indcamp/bilfait/boll_stat. English versions of these publications are available starting from 1998.

References

- Agarwal S, Pan J, Qian W (2014) Age of decision: Pension savings withdrawal and consumption and debt response. Working paper, National University of Singapore, Singapore, <http://ssrn.com/abstract=2487658>.
- Agarwal S, Qian W (2014a) Consumption and debt response to unanticipated income shocks: Evidence from a natural experiment in Singapore. *Amer. Econom. Rev.* 104:4205–4230.
- Agarwal S, Qian W (2014b) Access to housing equity and consumption: Evidence from a policy experiment in Singapore. Working paper, National University of Singapore, Singapore.
- Attanasio OP, Brugiavini A (2003) Social security and households' saving. *Quart. J. Econom.* 118:1075–1119.
- Attanasio OP, Rohwedder S (2003) Pension wealth and household saving: Evidence from pension reforms in the United Kingdom. *Amer. Econom. Rev.* 93:1499–1521.
- Bell B, Blundell R, Van Reenen J (1999) Getting the unemployed back to work: The role of targeted wage subsidies. *Internat. Tax Public Finance* 6:339–360.
- Beshears J, Choi J, Laibson D, Madrian BC (2011) Behavioral economics perspectives on public sector pension plans. *J. Pension Econom. Finance* 10:315–336.
- Beshears J, Choi J, Laibson D, Madrian BC (2012) The availability and utilization of 401(k) loans. Wise DA, ed. *Investigations in the Economics of Aging* (The University of Chicago Press, Chicago), 145–172.
- Blundell R, Pistaferri L, Preston I (2008) Consumption inequality and partial insurance. *Amer. Econom. Rev.* 98:1887–1192.
- Bottazzi R, Jappelli T, Padula M (2006) Retirement expectations, pension reforms, and their effect on private wealth accumulation. *J. Public Econom.* 90:2167–2213.
- Browning M, Crossley TF (2001) The life-cycle model of consumption and savings. *J. Econom. Perspect.* 15:3–22.
- Brugiavini A, Padula M (2003) Household saving behaviour and pension policies in Italy. Borsch-Supan A, ed. *Life-Cycle Savings and Public Policy* (Academic Press, New York), 99–146.
- Chetty R, Friedman JN, Saez E (2013) Using differences in knowledge across neighborhoods to uncover the impacts of the EITC on earnings. *Amer. Econom. Rev.* 103:2683–2721.
- Gale WG (1998) The effect of pension wealth on household wealth: A reevaluation of theory and evidence. *J. Political Econom.* 106:706–723.
- Jappelli T, Pistaferri L (2010a) Does consumption inequality track income inequality in Italy? *Rev. Econom. Dynam.* 13: 133–153.
- Jappelli T, Pistaferri L (2010b) The consumption response to income changes. *Annual Rev. Econom.* 2:479–506.
- Jappelli T, Pistaferri L (2014) Fiscal policy and MPC heterogeneity. *Amer. Econom. J.: Macroeconomics* 6:107–136.
- Jappelli T, Padula M, Pistaferri L (2008) A direct test of the buffer-stock model of saving. *J. Eur. Econom. Assoc.* 6:1186–1210.
- Lu TJ, Mitchell OS (2010) Borrowing from yourself: The determinants of 401(k) loan patterns. Working Paper 221, Michigan Retirement Research Center, University of Michigan, Ann Arbor.
- Meyer BD (2002) Labor supply at the extensive and intensive margins: The EITC, welfare, and hours worked. *Amer. Econom. Rev.* 92(2):373–379.
- Ministry of Economics and Finance (2014) Conto Annuale 2012 sul Pubblico Impiego, Rome. http://www.rgs.mef.gov.it/VERSIONE-I/Comunicazione/Note-per-la-stampa/2013/2013_conto_annuale/.
- Parker JA, Souleles NS, Johnson DS, McClelland R (2013) Consumer spending and the economic stimulus payments of 2008. *Amer. Econom. Rev.* 103(6):2530–2553.
- Shapiro MD, Slemrod J (1995) Consumer response to the timing of income: Evidence from a change in tax withholding. *Amer. Econom. Rev.* 85:274–283.
- Stephens M Jr (2001) The long-run consumption effects of earnings shocks. *Rev. Econom. Statist.* 83:28–36.