Istituto Nazionale Previdenza Sociale

ottobre 2024 – numero 83

WorkINPS *Papers*

A note on the effect of easing access to credit on firms employment

Federico Belotti
Edoardo Di Porto
Vadim Bottoni
Gian Marco Pinna¶

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.

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.

Tommaso Nannicini

A note on the effect of easing access to credit on firms employment

Federico Belotti
Department of Economics, Tor Vergata University of Rome

Edoardo Di Porto
University of Naples Federico II, CSEF and UCFS Uppsala University

Vadim Bottoni
National social security institute, INPS

Gian Marco Pinna
Department of Economics, Tor Vergata University of Rome

A note on the effect of easing access to credit on firms employment*

Federico Belotti†, Edoardo Di Porto‡, Vadim Bottoni§, Gian Marco Pinna¶

September 4, 2024

Abstract

This note investigates the effect of a policy measure implemented by the Italian government in 2014 called "Nuova Sabatini." This measure was aimed at easing access to credit for small and medium businesses, supporting investments in the acquisition of technological equipment. We exploit a difference-in-differences design to estimate the causal impact of the measure on different firm outcomes, namely capital stock, value-added, mean salary, and employment. Overall, we estimate that the measure significantly increased both firms' workforce and capital stock. Furthermore, we find very heterogeneous effects on employment by sector, size of the firm, and region of location. We then extend the analysis to include firms that select into treatment multiple times, finding evidence that granting access to credit multiple times enhances the effectiveness of the measure compared to firms applying just once.

Keywords: Industry 4.0, employees, easing credit, staggered difference in differences, multiple treatment

JEL: H32, D22, D24, J00

*We thank Fabio Pagliarini and his team at MISE for providing the data to Inps, the data used in this note were managed by Inps personnel only, the external authors provided theoretical and empirical insights, programs and ideas to develop the final version of this note.
† Department of Economics, Tor Vergata University of Rome - Email: federico.belotti@uniroma2.it; ‡ University of Naples Federico II, CSEF and UCFS Uppsala University - Email: edoardo.diporto@unina.it; § National social security institute, INPS - Email: vadim.bottoni@inps.it; ¶ Department of Economics, Tor Vergata University of Rome - Email: gianmarco.pinna@uniroma2.it.
The views expressed in the paper are those of the authors only and do not involve the responsibility of the Bank of Italy, the Eurosystem, or the Italian Social Security Institute (INPS).
This project was realized thanks to the VisitINPS Scholars programme.

Abstract

Questa nota si propone di indagare gli effetti di una politica adottata dal governo italiano nel 2014, chiamata "Nuova Sabatini". Lo scopo di questa misura era di rendere più semplice l'accesso al credito per le piccole e medie imprese, in modo tale da favorire gli investimenti nell'acquisizione di beni di tipo tecnologico. In questo lavoro adottiamo la metodologia delle difference-in-differences per stimare l'effetto causale della misura su diversi outcome d'impresa, come, ad esempio, lo stock di capitale, il valore aggiunto, il salario medio e l'occupazione. Sommariamente, troviamo che la misura abbia aumentato in maniera significativa la forza lavoro e lo stock di capitale delle imprese che ne hanno fatto uso. Inoltre, troviamo effetti eterogenei sull'occupazione per settore, dimensione e collocazione regionale dell'impresa. Infine, estendendo l'analisi con l'inclusione delle imprese che hanno fatto uso della misura per più di una volta, troviamo evidenze empiriche che garantire alle imprese di poter usufruire della misura anche più di una volta ne aumenta gli effetti positivi rispetto al caso delle imprese che hanno adottato la misura una sola volta.

1 Introduction

The Fourth Industrial Revolution (4IR) is a global trend towards automation and data exchange in manufacturing technologies and processes, which is inducing profound societal and industrial transformations. Over the recent years, many countries have implemented policy measures to foster the adoption of Industry 4.0 technologies. A notable example is Germany's "Industrie 4" program, which was a policy project that advocated for digitization, organizational evolution, and productivity enhancement in the German manufacturing system. This initiative spurred similar policies globally, including the one analyzed in this study, the "Nuova Sabatini" (NSAB), a measure provided by the Italian Government aimed at easing access to credit for small and medium businesses supporting investments in the acquisition of equipment.[1]

The motivation for this analysis stems directly from the rising popularity that industrial policies have been enjoying in recent years among policymakers, which naturally calls for economists to deliver thorough and reliable assessments of the effectiveness and efficiency of these policies. However, besides efficacy and efficiency considerations, it is equally important to investigate whether these policies inadvertently lead to adverse outcomes beyond their intended objectives, as pointed out by Juhász et al. [2023]. A crucial aspect of concern for policies that enhance capital investments is, in fact, their potential indirect effects on the labor market.

The identification of industrial policies' causal effects can pose substantial challenges. Measures like the NSAB are usually implemented through a national call, and as a result of the administrative process, firms select to participate in the program. Moreover, when these programs last for years and are implemented in subsequent waves, firms may self-select multiple times.

This paper addresses these issues by providing a credible evaluation of the NSAB's causal effects for firms in different application cohorts, requesting NSAB one or multiple times. Our analysis unfolds in the following manner. First, we assess whether the policy achieves its intended goals by estimating its effects on the firms' capital stock and value-added. Second, we investigate the impact of NSAB on workers' employment and mean salary. While the effect on capital is expected to be non-negative, the impact on employment is ex-ante unclear since it depends on the degree of complementarity (or substitutability) of labor with the newly acquired equipment. Third, we dig into the heterogeneity of these effects to provide insightful recommendations on fine-tuning the requirements to access the policy measure. Fourth, while the preceding analysis focuses on firms that apply for NSAB only once over multiple waves, we extend the investigation to include firms making multiple applications.

[1] Equipment, such as machinery, productive assets, hardware, as well as software and digital technologies (Legislative Decree n.69, June 21, 2013).

Our empirical strategy exploits a difference-in-differences design. To deal with possible heterogeneous treatment effects in various cohorts and across various firms' dimensions (i.e., size, age, location) and staggering treatment, we employ the estimator proposed by Callaway and Sant'Anna [2021]. In our context, firms that have not yet applied for the NSAB in a particular wave serve as the control group. Therefore, our identification strategy assumes parallel trends between "not-yet" and treated firms within a specific wave.

Our main findings indicate that easing credit access raises capital and employment significantly and persistently. We find that the average treatment effect on firms with access to NSAB is equal to 0.5 full-time equivalent (FTE) employees in the first year. After five years, the effect is six times larger. Interestingly, we do not observe significant effects on average earnings and value-added.

We find significant heterogeneity in treatment effects by sector, firm size, and region of location. Manufacturing, small and micro firms located in the North of Italy benefited the most from the policy. On the other hand, firms operating in the construction sector or that are located in the South of Italy seem not to be significantly impacted by the measure. Finally, we find cases where employment is negatively affected, such as medium-sized firms and firms in the logistics and transport sectors, suggesting a certain degree of substitutability between newly acquired capital and labor in these types of firms. Policymakers aiming to mitigate job loss could exploit this evidence by targeting the policy toward firms showing a higher degree of complementarity between capital and labor, thus excluding firms in logistics.

In our final evaluation exercise, we investigate the effects of the treatment for firms applying for NSAB more than once. Indeed, policymakers might find it helpful to know whether endowing firms with further financial support brings increasing or diminishing returns compared to granting it just once. We provide evidence to answer this question by considering firms applying twice. As a pre-treatment period, we consider the time between the first and the second application. In this case, firms that have applied once for the NSAB in a particular wave but have not yet applied for the second time serve as the control group, i.e., the "not-yet-two-times" treated firms. Compared to firms accessing NSAB only once, endowing firms with further financial support brings about 0.7 additional FTE to the mean recipient firm in the first year. We interpret this result as supportive evidence of an extended support for investments in technological equipment.

The closest study to our work would be Cingano et al. [2022]. Unlike our case, the authors employed a regression discontinuity design to estimate the effects of public investment subsidies granted to Italian firms operating in disadvantaged areas. They find that such subsidies raised investments by 39% and employment by 17% over six years in firms near the cutoff (i.e., firms that narrowly won

the competition for the subsidy compared to those that narrowly lost it). Differently from the policy analyzed in Cingano et al. [2022], NSAB is granted to any firm that complies with the eligibility criteria delineated by the law.[2] This posed additional problems from a methodological point of view, which we addressed by using a staggered difference-in-differences design, as outlined earlier.

The rest of this paper is structured as follows. In section 2, we present the data used to carry out the analyses. Section 3 outlines the empirical strategy and the reasons behind its choice. Section 4 presents the results and digs into the heterogeneity of the effects across different groups of treated firms. Section 5 report the analysis on firms treated twice, while Section 6 draws some conclusive remarks.

2 Data

To conduct the analysis, we employed two primary data sources. The first comes from MISE (the Italian Ministry of Economic Development), from which we collect information on firms' planned investments tied to the requested subsidy. More precisely, this dataset contains information on the amount of the financial support, the type (whether it is classified as "ordinary" or "technological"), the date on which the application was submitted, and its date of approval.[3] The second data source is the Italian Social Security Institute (INPS), from which we retrieve all the relevant firms' characteristics, such as the number of employees, the total payroll, firms' age, and sector. In addition, we complement these datasets with balance sheet information extracted from Cerved[4] to carry out the part of the analysis relative to value-added and capital.

The sample of interest encompasses the firms that demanded and subsequently received the subsidy between 2014 and 2019. In order to be eligible, firms had to comply with some criteria: first, they had to be classified as small-to-medium enterprises (SMEs), meaning that they could not have more than 250 employees; second, firms could not be experiencing financial distress (i.e., they could not be in the process of closing or under an insolvency procedure); third, they needed to have their headquarters in Italy by the end of the planned investment period; fourth, firms could not be operating in the insurance and financial sectors, or having activities linked to exports. The original MISE dataset includes about 120,000 requests for the subsidy over the 2014-2019 period. Out of these firms, we select only those that got their requests approved, which trims the sample to about 36 thousand firms.

[2] Firms should: i) be small or medium enterprises (maximum 249 employees); ii) be healthy; iii) operate in Italy in any economic sector except for those linked to exports or the financial or insurance sector.

[3] "Ordinary" investments were subsidized at a slightly inferior rate compared to "technological" types. The size of the subsidy for the former type of investment was equal to 2.75% of the total amount of the planned investment, whereas the latter types of investments received a subsidy equal to 3.575%.

[4] Cerved is a database that includes all limited liability companies incorporated in Italy.

Furthermore, firms could request the subsidy more than once. Thus, the total number of firms that requested the subsidy is lower than the number of subsidies granted. For our main specification, we initially discard firms that received more than one subsidy. After this exclusion, the resulting sample is an unbalanced panel of 22,200 firms observed over 2010-2019, for a total of 198,618 observations. Moreover, it is important to stress that the size of each cohort (the set of firms that apply for NSAB in any given year) significantly varies over time. This may be due to the fact that the financial support for "technological" types of investments was only introduced in 2017, which coincides with the number of accepted requests more than doubling in that year compared to the previous one (see Table A1 in Appendix A). In addition, the support granted for the technological types of investment is slightly more generous than the one for "ordinary" types, which might have incentivized more firms to apply.[5] The composition of firms in the dataset also varies greatly across economic sectors: about 54.5% of the firms that received the subsidy operate in the manufacturing sector, another 35% are, more or less, evenly split across the construction, the wholesale and retail trade, and the transportation and storage sector, while the remaining 10.5% is spread across the remaining sectors (see Table A2 in Appendix A). In terms of firm size (measured in the year before receiving the subsidy), micro (<10 employees) and small (more than nine and less than fifty employees) firms together constituted more than 90% of the sample (see Figure A4 in Appendix A). The rest of the firms in the sample are medium-sized (more than 49 and less than 250 employees). Finally, the median firm in the sample was 18 years old, employed 11 employees, and paid about 20 thousand euros per year per employee (see Table A3 in Appendix A).

3 Empirical strategy

The recent literature on the estimation of heterogeneous treatment effects in DiD designs with variation in treatment timing (staggered adoption) shows the drawbacks associated with using standard TWFE linear regression specifications [Borusyak et al., 2024, De Chaisemartin and d'Haultfoeuille, 2020, Goodman-Bacon, 2021, Sun and Abraham, 2021].

In our setting, firms request financial incentives set up by the NSAB over different annual waves. Due to the diverse production structures and other characteristics of these applicants, we expect that easing access to credit may have heterogeneous treatment effects across different types of firms. Moreover, the scope of the NSAB policy widened over time. Between 2014 and 2016, the policy supported only "ordinary" types of investments, while from 2017 to 2019, "technological" types of

[5] "Ordinary" types of investments were subsidized at 2.75%, whereas "technological" types of investments received a support equal to 3.575% of the planned investment.

investments were also included in the program. The practical implication of this dynamic is that later cohorts of applicants became larger in size and received, on average, slightly more generous support.[6] While this difference is not enough to call for a separate analysis, we worry that it could partly result in heterogeneous treatment effects across cohorts. Furthermore, while the last cohort in our sample could be considered a good comparison group for the second-to-last treated cohort (as both the sample size and the mixture of the two types of subsidies are similar), it is less the case for earlier-treated cohorts. All these elements justify our choice of the estimator.

Our analysis exploits the estimator proposed in Callaway and Sant'Anna [2021, CS]. The reason for this choice is threefold: first, the CS approach does not impose any restriction on treatment effect heterogeneity across groups and/or across time. Second, it embeds different aggregation schemes that can be used to highlight how average treatment effects vary across different dimensions; third, this methodology exploits "naturally" not-yet-treated units as a control group to the fullest possible extent.

To get our main results, we run the analysis without including any covariates, as pre-trend coefficients suggest that the parallel trends assumption holds even unconditionally. One small caveat to remember when reading the results is that both the treatment and the control group are affected by compositional changes. Besides this caveat, the interpretation of the results is fairly straightforward. Our estimates represent the average treatment effect on the treated on capital, value-added, employment, and wages across all treated cohorts receiving the NSAB support.

4 Empirical Results

In our baseline specification, we are interested in measuring the causal effects of the financial help provided by the NSAB on various outcomes of the beneficiary firms. We start first by presenting the findings on the impact on employment and average salary in recipient firms.

Figure A shows the event-study coefficients: green dots display pre-treatment ones, whereas red dots display post-treatment ones. Pre-treatment coefficients suggest the absence of a pre-trend and that the parallel trend assumption holds. We find an average treatment effect of about 0.46 in the year in which firms receive the subsidy, which in the fifth year after the treatment becomes 3.1. The subsidy seems to stimulate recipient firms to hire about a part-time employee in the first year after the adoption of the policy and three full-time employees after five years.[7]

[6] Even though the difference between technological and ordinary types of investments amounts to less than a percentage point of the total planned investment (0.825%), technological investments are subsidized 30% more compared to ordinary types.

[7] However, it should be remarked that when interpreting these results, one needs to be aware that the dynamic estimates shown in the chart are affected by compositional changes. For an in-depth review of the issue, see Callaway and Sant'Anna [2021].

Figure A: Total variation in the number of employees

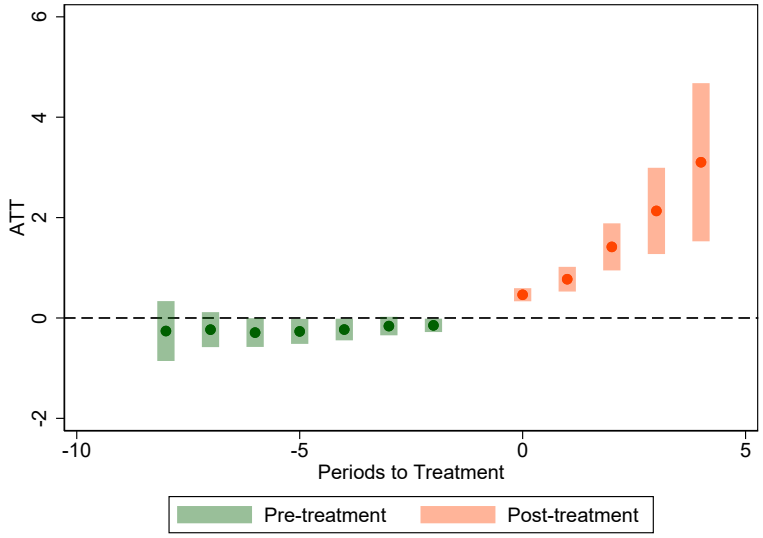

Subsequently, we analyzed the policy's impact on the average salary of the beneficiary firms. Figure B presents estimates for this outcome. We observe that the policy does not have a significant impact on average real wages.

Figure B: Total variation in employees' mean salary

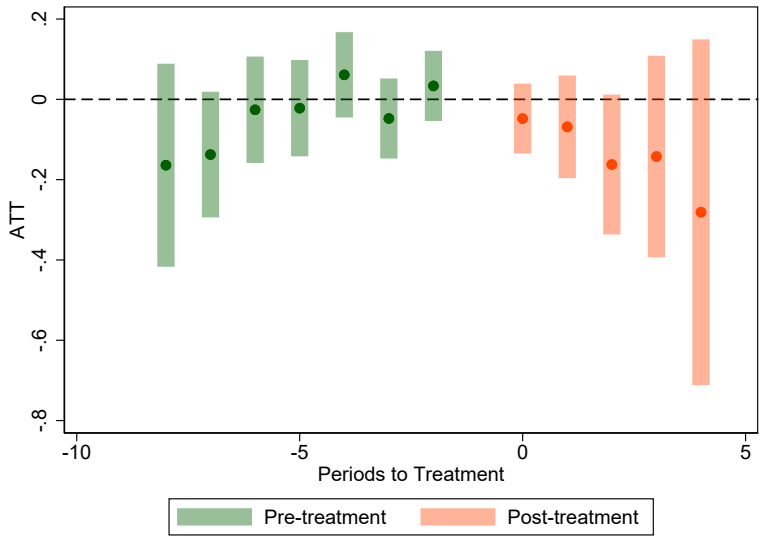

Finally, we looked at the effects of the policy on capital and value-added. The results are provided in Figure C. We observe that post-treatment coefficients are, in general, positive and significant (with the exception of $e = 0$ for value-added). However, while pre-treatment coefficients for the capital stock outcome suggest that the parallel trend assumption behind our identification holds, this is not verified for the case of the value-added specification. Thus, while there is enough evidence to argue a positive

and long-lasting effect of the policy on capital, we do not have enough evidence to support the same claim for the value-added dimension.

Figure C: Average treatment effect on capital and value-added

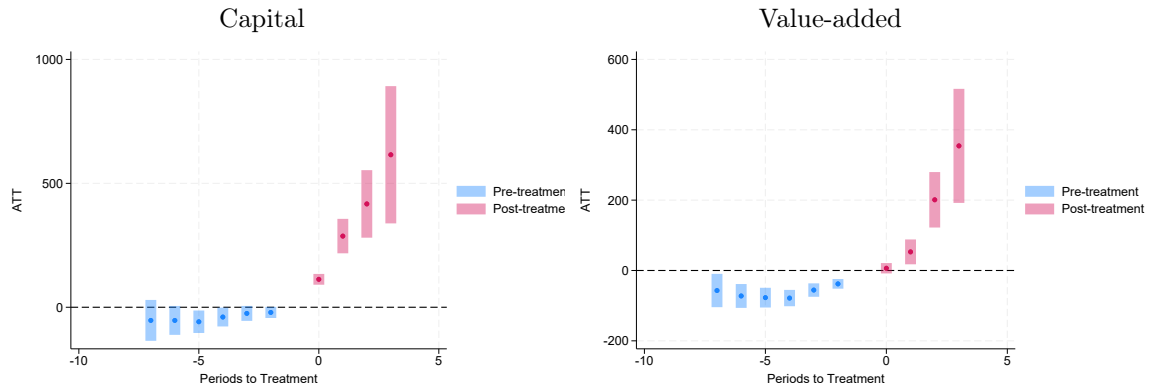

4.1 Heterogeneous impacts

In this section, we will try to provide an answer to the following question: Did the policy measure have heterogeneous effects on employment across firms with different characteristics? To answer this question, we run separate regressions, splitting the sample according to various observable firm characteristics.

First, by looking at the sectoral decomposition of firms (see Table A2 in Appendix A), we noticed that about 90% of them work in only four sectors: manufacturing, construction, wholesale and retail trade, transportation and storage sector. Therefore, we split the sample and repeated the analysis for these four sectors separately, plus the aggregation of the remaining sectors. Figure D below displays the results.

We observe that the manufacturing sector's chart closely resembles the results we found in our main specification, reported in Figure A. Coefficients are only slightly larger with respect to our baseline specification. On the other hand, the charts for the other three sectors, plus the aggregation of the other ones, look rather different. Estimates for the construction sector are close to zero and not statistically significant. For the wholesale and retail sector, we do find instead three positive and significant at the 10% level estimates. Strikingly, we even observe one negative and significant estimate for the transportation and storage sector in the post-treatment period. However, we also notice that pre-treatment coefficients might suggest a violation of the parallel trend assumption in this case. Finally, the chart showing the results for the set of the remaining sectors also exhibits significant and positive coefficients for the first four periods after the treatment. Overall, figure D showcases a remarkable

sectoral heterogeneity of treatment effects, suggesting that our baseline results are mainly driven by firms working in the manufacturing sector, which, in fact, represents more than 50% of the sample.

Figure D: Sectoral heterogeneity in total variation in the number of employees

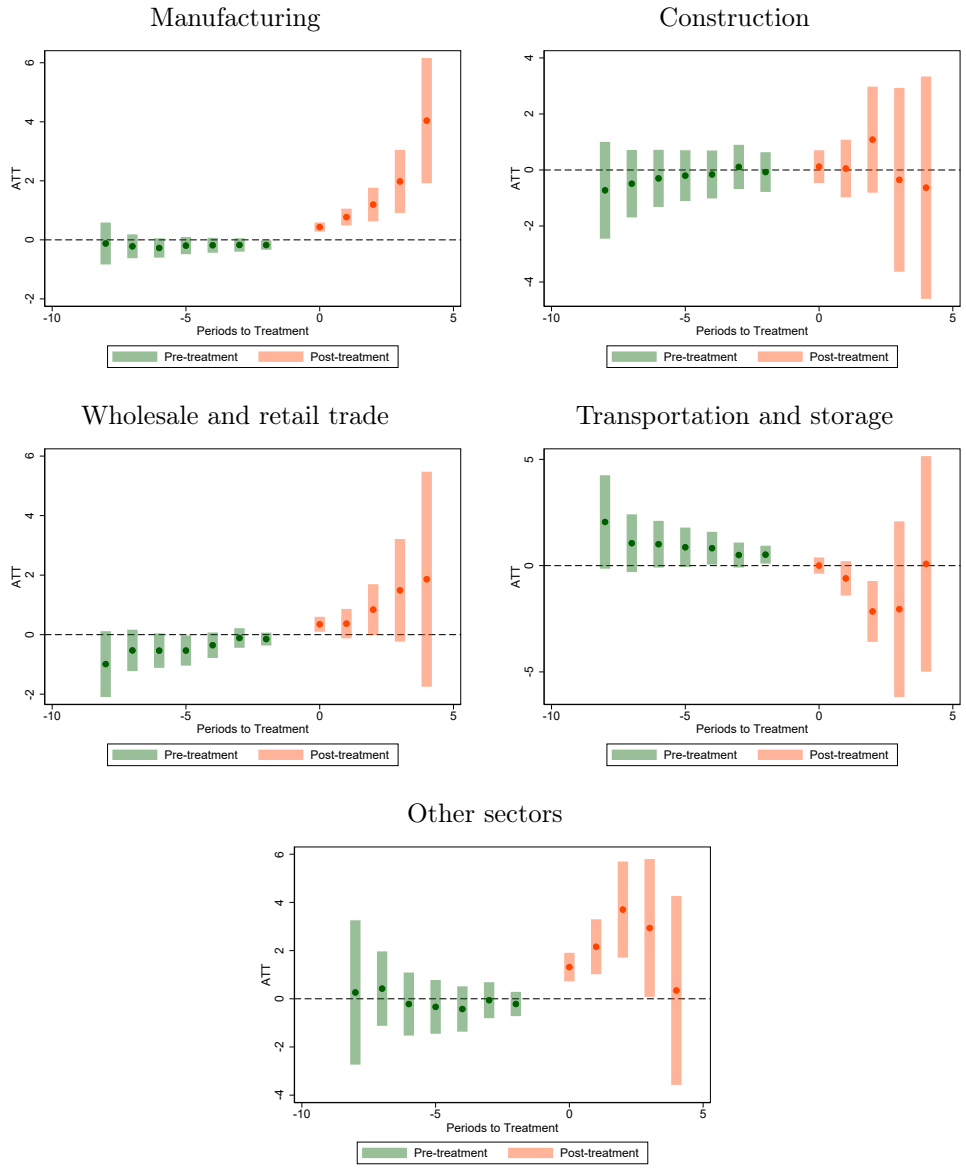

Secondly, we checked for possible heterogeneous effects in terms of firm size. To do that, we divided the sample into micro, small, and medium enterprises according to each firm's number of employees in the year before they accessed the NSAB. Figure E shows the results.

The effect of the policy on firms' number of employees is significant, positive, and persisting only for micro and small firms. On the contrary, employment in medium firms seems to be impacted negatively by the subsidy. These results hint that helping firms acquire new capital may have different implications according to their size. A potential explanation for these findings could be that, while for

micro and small firms, subsidizing the acquisition of new capital may represent a vehicle to overcome the typical obstacles that these firms encounter when they try to expand, for medium-sized firms, instead, acquiring new capital serves the purpose of transforming their production processes to make them more flexible and efficient. This means that the degree of substitutability between capital and labor typically varies according to the firm's size.

Figure E: Heterogeneity in terms of firm size in total variation in the number of employees

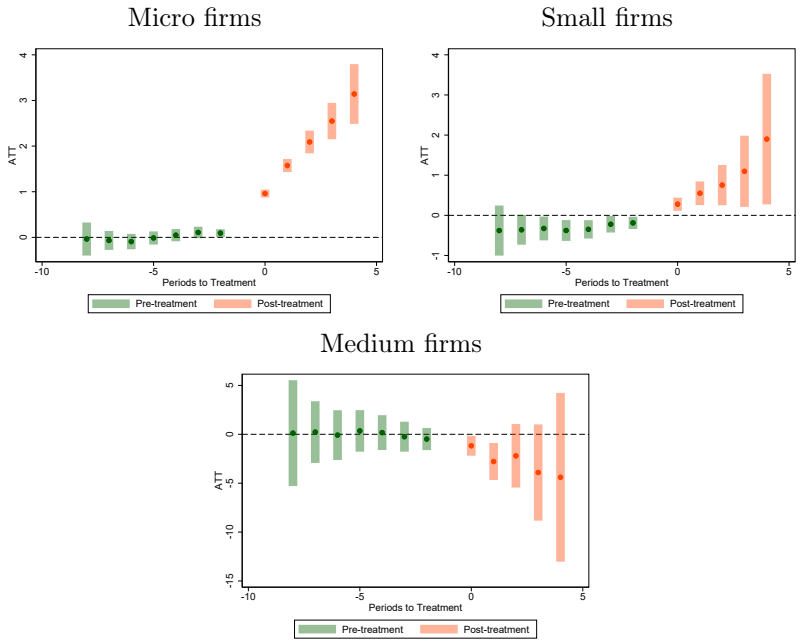

Finally, we evaluate the effects of the subsidy in different macro areas, grouping regions in the North, the Center, and the South of Italy. Since the South of Italy is industrially less developed than the North, if the subsidy allows firms to overcome financial barriers to growth, one would expect to observe a larger impact in the South rather than the North of Italy. However, the findings contradict this hypothesis.

Figure F below shows the regional heterogeneity in the treatment effects. The chart that displays estimates for the North of Italy is, again, remarkably similar to the one obtained using the whole sample. As a matter of fact, a striking 72% of the policy's recipients were firms located in the North. Looking at the top-right picture, showing estimates for the Center, we notice instead that pre-treatment coefficients exhibit an ascending trend, which indicates that treated firms may have already been on a comparatively positive growth trend before the treatment kicked in. In this case, the parallel trend assumption does not seem to hold. Nevertheless, we also notice positive and significant coefficients. On the other hand, the chart for the South does not show significant effects of the NSAB on employment for the post-treatment period. While these results may seem counterintuitive, it may be the case

that the economic and industrial complex that is in place in the South requires a more structural and comprehensive intervention to activate firms' growth with respect to the one that is in place in the North of Italy.

Figure F: Regional heterogeneity in total variation in the number of employees

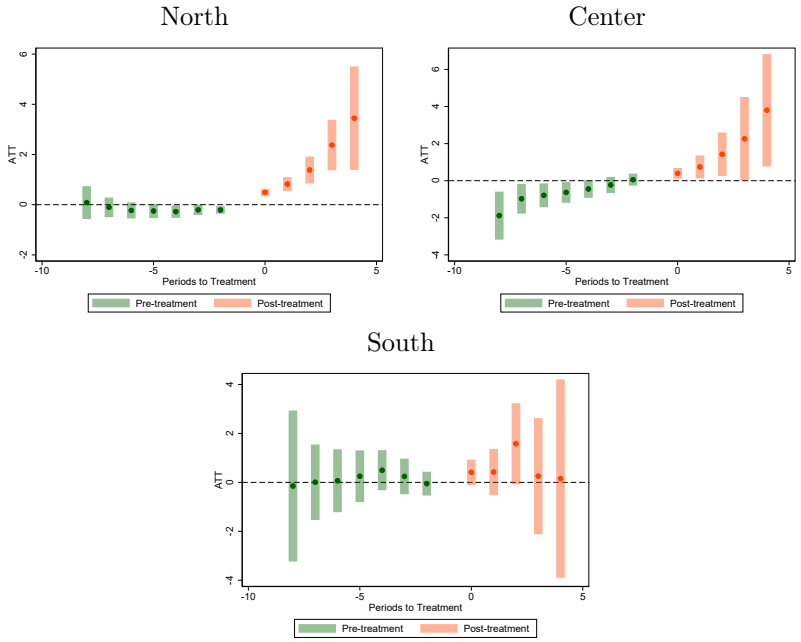

4.2 Multiple treatment

So far, we have conducted the analysis by excluding those firms that have received more than one "dose" of the treatment. To be precise, we excluded from the sample firms that received more than one subsidy within any given year or across the entire considered time frame. As an addendum to our main analysis, we now delve into the estimation of the average treatment effect for firms that received more than one subsidy, as long as these subsidies were spread across separate years. Estimating the effect of a further "dose of treatment" has important policy implications. Indeed, policymakers might find it relevant to have an answer to the following question: does a second dose of treatment have a magnifying or diminishing return with respect to the first one? Supposing the latter case turns out to be true, policymakers might, in fact, deem it more efficient to divert these resources that would have been assigned to firms that already received the subsidy elsewhere (for instance, to firms that have not received the subsidy yet).

However, estimating the effect of multiple treatments (or of different treatment intensities) is a

tricky endeavor, implying being willing to strengthen the parallel trend assumption. Some recent works have tried to provide the theoretical foundations and some practical solutions to this econometric problem (De Chaisemartin and d'Haultfoeuille, 2018; Callaway et al., 2021; De Chaisemartin and D'haultfœuille, 2023). The closest one to what is going to be presented next is Fricke [2017]. The object of analysis of this paper is the case in which researchers would like to estimate the average treatment effect of one dose of treatment compared to another dose of lower intensity. In this particular scenario, the key assumption to be made is that units that received the higher-intensity treatment would have responded to the lower-intensity one in the same way as units that actually received the lower-intensity treatment (or vice versa, if the comparison group is the one of higher-intensity-treated units). In other words, making this comparison implies strengthening the parallel trend assumption by imposing homogeneity of treatment effects on one of the two groups of treatment units. However, this assumption might sometimes be too strong, especially in cases in which the treatment is not randomly assigned. Thus, the author of this paper proposes an alternative interpretation for the DiD estimand for cases when it is hard to assume that this kind of homogeneity holds. Rather than the difference between the two treatments, the DiD estimand can be then interpreted in these cases as the lower bound of an average treatment effect compared to the case of no treatment.

The case we are analyzing, though, is a simplified version of this problem, as we aim to estimate the effect of a further treatment dose of the same proportion as the first one. To do so, we simply need to retain in the sample those firms that are treated at least, and at most, twice and discard the rest. Then, we also get rid of the years prior to the first treatment and replicate the estimation in the same fashion as in our main analysis, using the years between the first and the second treatment as "pre-treatment" periods.[8] In this simplified case, we do not need to make any assumption on the homogeneity of treatment effects between firms that received the treatment only once and those that received it twice because we can actually measure these effects.

In our case, we feel confident in presenting results on units treated up to two times. However, the same procedure could be applied to units treated as many times as one wishes, as long as the sample does not shrink so much that it affects the credibility of the resulting estimates.

Figure G illustrates the results of this exercise. As explained in the previous paragraph, our control group is composed of units that have accessed the NSAB once but have yet to receive the treatment a

[8]Notice that an alternative way to proceed would have been to use firms only treated once as a comparison group for firms treated twice. However, this kind of comparison implies a stronger parallel trend assumption with respect to the framework outlined above. More specifically, we would need to assume the absence of unobservable reasons that may induce firms to select into the group of firms treated only once or of firms treated twice. In other words, we would have to argue that there are no fundamental hidden features between these two groups of firms that strongly predict their sorting into one of the two groups. Since we can avoid making this assumption, we prefer to stick with the setup that exploits later-treated units as a comparison group.

second time. The shown estimates can thus be interpreted as the average treatment effect of a second dose of treatment with respect to the first dose.

Figure G: Average treatment effect on the treated of the second treatment for firms that received two subsidies

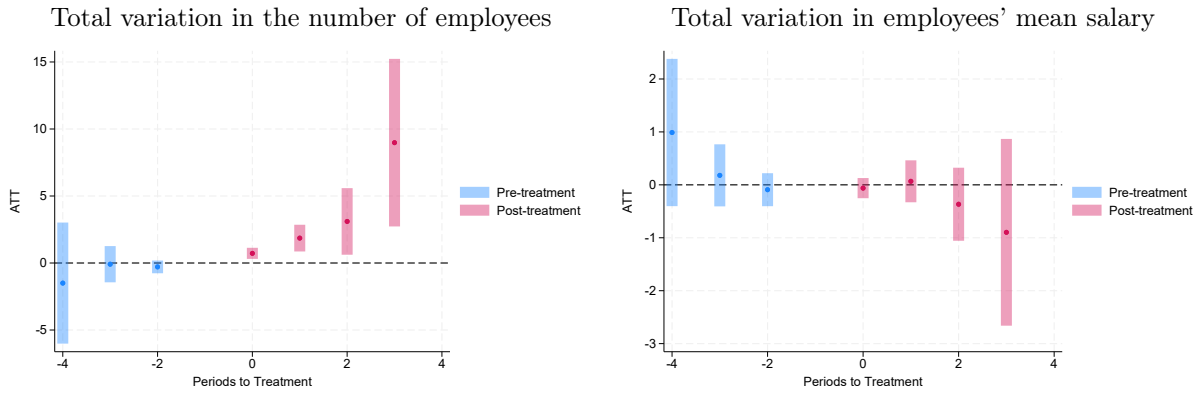

As in our main analysis, we do not find any significant effect on mean wages. On the other hand, the effect on the total employees' growth is positive, large, and significant. On impact, obtaining the subsidy again brings 0.72 additional employees to the mean recipient firm. After four years, firms have added, on average, nine additional employees with respect to firms that have received the subsidy only once. From a policy evaluation point of view, the evidence we find points to very large returns from granting the subsidy to firms that have already obtained it once.

5 Conclusions

In this note, we studied the effects of a policy aimed at creating incentives for firms to innovate. These incentives took the form of financial support to firms that desired to undertake technologically advancing investments to modernize their capital equipment. We analyzed the effect of this policy on capital, value-added, mean salary, and employment.

We found evidence that the policy had a significant, positive, and long-lasting effect on capital. The evidence is, instead, more blurred for what it concerns value-added. On the contrary, we do not find any evidence that the policy affected the average salary in recipient firms. More importantly, the main focus of our analysis was to study the effects of these subsidies on employment. While the impact of the policy on capital, at least on the year of the treatment, is almost trivial (given that the subsidy was granted only conditional on the fact that the receiving firm could prove that it carried out the planned investments), the effects of the subsidy on employment is, instead, ex-ante unclear. The net

outcome of the policy on occupation, in fact, crucially depends on the degree of complementarity (or substitutability) of the acquired capital and labor.

We found that, overall, the policy had a positive and persistent effect on employment. Nevertheless, this result was not homogeneous for all types of firms. Investigating possible heterogeneous treatment effects across various dimensions, we discovered that the overall growth in employment was mostly happening in micro and small firms (while it decreased in medium firms) located mainly in the North of Italy and operating in the manufacturing sector. Finally, we evaluated the effect of a second subsidy with respect to the first one and found that firms that got the subsidy twice were able to grow their workforce faster than firms that got the subsidy only once.

These results have critical implications for the design of industrial policy. First, they suggest that for micro and small firms, capital is not a substitute for labor. Instead, helping small-sized firms grow their capital stock can actually enable them to expand their workforce. However, this is not true for medium-sized firms, where capital and labor seem to have some degree of substitutability. This hints at the existence of different elasticities between capital and labor according to the firm's size. Secondly, subsidies worked mainly in the North, where prevalent economic structures are already strong, facilitating industrial growth. Firms in the South of Italy are, instead, both underrepresented in the sample and seemed to have benefitted less from the policy. Lastly, our results suggest that allowing firms to exploit the policy's benefits more than once brings largely positive outcomes to recipient firms.

References

Kirill Borusyak, Xavier Jaravel, and Jann Spiess. Revisiting Event-Study Designs: Robust and Efficient Estimation. *The Review of Economic Studies*, 02 2024.

Brantly Callaway and Pedro HC Sant'Anna. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230, 2021.

Brantly Callaway, Andrew Goodman-Bacon, and Pedro HC Sant'Anna. Difference-in-differences with a continuous treatment. *arXiv preprint arXiv:2107.02637*, 2021.

Federico Cingano, Filippo Palomba, Paolo Pinotti, and Enrico Rettore. Making subsidies work: rules vs. discretion. *Bank of Italy Temi di Discussione (Working Paper) No*, 1364, 2022.

Clément De Chaisemartin and Xavier d'Haultfoeuille. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–96, 2020.

Clément De Chaisemartin and Xavier d'Haultfoeuille. Fuzzy differences-in-differences. *The Review of Economic Studies*, 85(2):999–1028, 2018.

Clement De Chaisemartin and Xavier D'haultfœuille. Two-way fixed effects and differences-in-differences estimators with several treatments. *Journal of Econometrics*, 236(2):105480, 2023.

Hans Fricke. Identification based on difference-in-differences approaches with multiple treatments. *Oxford Bulletin of Economics and Statistics*, 79(3):426–433, 2017.

Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 2021.

Réka Juhász, Nathan J Lane, and Dani Rodrik. The new economics of industrial policy. Technical report, National Bureau of Economic Research, 2023.

Liyang Sun and Sarah Abraham. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199, 2021.

A Appendix

Table A1. Total number of subsidies

Year	Number of firms	Percentage
2014	826	3.72
2015	1067	4.81
2016	2495	11.24
2017	6009	27.07
2018	6523	29.38
2019	5280	23.78
Total	22200	100

Notes: figures refer to firms that received treatment once.

Table A2. Distribution of subsidies across sectors

Sectors	Number of firms	Percentage
C	12100	54.5
F	2082	9.38
G	2666	12.01
H	3013	13.57
Others	2339	10.54

Notes: figures refer to firms that received treatment once. Sector C is manufacturing, sector F is construction sector G is wholesale and retail trade, sector H is transportation and storage.

Table A3. Age, avg number of employees, mean wage

Stats	Firms' average age	Total number of employees	Annual mean wage
Mean	19,8	18.6	20185
Median	18	11	20156

Notes: figures refer to firms that received treatment once. Mean wages are rounded to the closest integer.

Figure A4: Size of firms that received subsidies as a % of the total

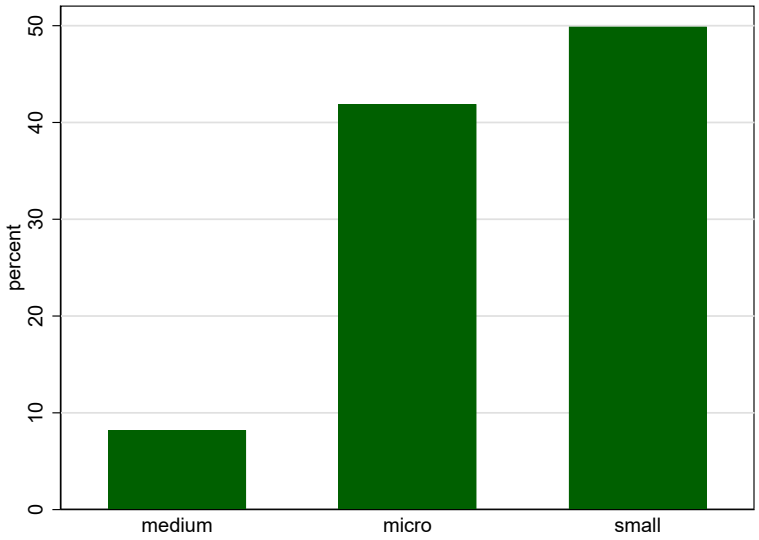

Source: MISE data, authors' calculations.
Notes: figures refer to the sample of firms treated only once. Micro firms have <10 employees, small firms more than 9 and less than 50 employees, medium firms have more than 49 and less than 250 employees.

Figure A5: Location of firms that received subsidies as a % of the total

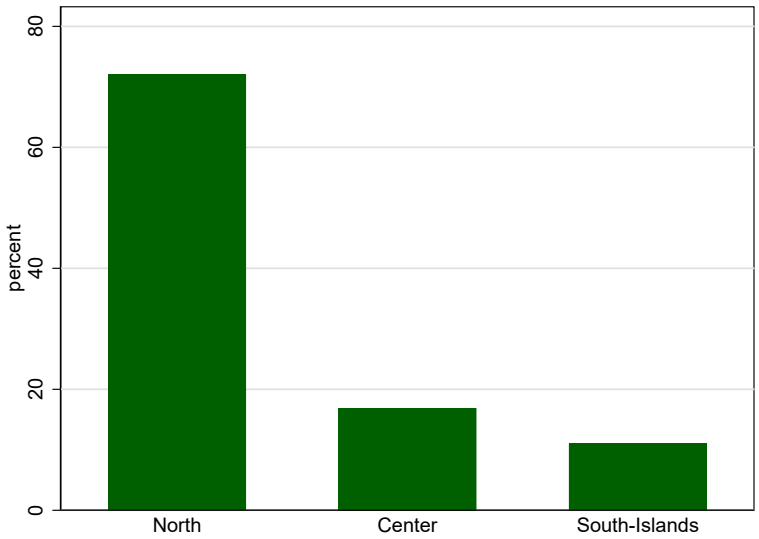

Source: MISE data, authors' calculations.
Notes: figures refer to the sample of firms treated only once. Regions in the North of Italy include: Valle d'Aosta, Veneto, Trentino-Alto Adige, Lombardia, Liguria, Friuli-Venezia Giulia, Emilia-Romagna, and Piemonte. Regions in the Center are: Lazio, Marche, Toscana, and Umbria. Regions classified as South or Islands are: Abruzzo, Basilicata, Calabria, Campania, Molise, Puglia, Sardegna, Sicilia.