Entrepreneurs and Risky Spouses*

ISAAC HACAMO† and VALENTINA RUTIGLIANO§

March 27, 2026

PRELIMINARY

Please do not circulate

ABSTRACT

Most entrepreneurs live in households with multiple members, yet the existing academic work treats them as isolated individuals, overlooking how intra-household risk sharing may shape firm-level decisions. Using an exogenous increase in labor income risk, we show that entrepreneurs married to exposed workers reduce firm size and leverage, and increase their demand for liquid assets. Additional empirical evidence, using proxies for risk aversion and firm riskiness, supports a precautionary savings motive. We develop a macro model with risk-averse entrepreneurs who have heterogeneous productivity and face borrowing constraints. Relative to a conventional economy with entrepreneurs and workers as separate agents, an economy featuring entrepreneurs married to workers leads to greater concentration of capital among the most productive entrepreneurs. Following a negative productivity shock, concentration further increases, but total output declines relative to a conventional economy.

JEL classification: D13; G51; G32; L26

Keywords: Intra-household transmission; household finance; precautionary savings; entrepreneurship; unemployment risk

*The views expressed herein are those of the authors and do not necessarily reflect the views of Statistics Canada. We thank seminar and conference participants at Deakin University, Federal Reserve Bank of Chicago, UNSW Sydney, and Southern Methodist University. We also thank Gene Amromin, Mariacristina De Nardi, Xiaoyang Li, Gordon Phillips, Eddie Shore, and Wing Wah Tham for very helpful comments.

†Kelley School of Business, Indiana University. Email: ihacamo@iu.edu.

§Cox School of Business, Southern Methodist University. Email: vrutigliano@smu.edu.

1 Introduction

Over 70% of entrepreneurs in the U.S. are married, and approximately 60% have children.[1] Yet, existing academic work rarely considers entrepreneurs explicitly as members of multi-person households. Considering entrepreneurs as part of households implies that they are exposed to uninsurable household-level shocks. These shocks may shape firms' investment choices: entrepreneurs may seek to mitigate household risk by adjusting investment and firms' risk exposure. These actions may, in turn, feed back into the broader economy, propagating and amplifying the shocks that give rise to household risk. Despite the potential importance of this *household-entrepreneur channel*, our understanding of its implications remains limited.

Whether household income risk should affect entrepreneurial firms is not theoretically obvious. If entrepreneurs have low risk-aversion or do not face borrowing constraints, such shocks should not change firm investment or risk exposure. In contrast, if entrepreneurs have high risk-aversion and family income is used to relax borrowing constraints, entrepreneurs may adjust investment and risk exposure in equilibrium. In a simple framework with risk-averse entrepreneurs facing borrowing constraints tied to pledgeable family income, income shocks lead to tighter borrowing constraints that reduce firm size and leverage. Precautionary saving motives further strengthen these effects as demand for safer and liquid assets increases.

We find that entrepreneurs whose spouses are exposed to a rise in labor income risk reduce investment and leverage. Further, consistent with precautionary motives, these entrepreneurs increase their demand for liquid assets as they increase dividend payouts. These effects are stronger for older and female entrepreneurs—especially when they have greater bargaining power within the household—and also for firms with more volatile cash-flows. These findings highlight the importance of intra-household risk sharing when studying firm-level decisions of entrepreneurs. We then develop a macro model with heterogeneous entrepreneurial productivity to examine the implications for wealth concentration and total output.

To identify the effect of household income risk on entrepreneurial firms, we exploit the sharp and largely unanticipated collapse in global oil prices in 2014 as a source of increased

[1] See, e.g., https://www.kauffman.org/reports/the-anatomy-of-an-entrepreneur/.

income risk for workers in the oil industry. The shock was driven by global supply and demand forces, as well as an unexpected decision by OPEC to maintain production despite falling prices, and is widely regarded as exogenous from the perspective of Canadian firms and households (e.g., Arezki and Blanchard, 2015; Kilian and Baumeister, 2016). This unexpected shock allows us to construct treatment and control groups, as well as a placebo test that significantly strengthens our identification strategy.

We define treated entrepreneurs as those whose spouses were employed in the oil sector. The treatment group consists of entrepreneurs whose spouses were working in the oil sector in 2013–2014, prior to the price collapse, and were therefore unexpectedly exposed to the shock. We then construct a control group by matching each treated entrepreneur to an entrepreneur in the same four-digit industry, gender, and commuting zone, whose spouse has similar income but works outside the oil sector. Our design compares firms operating in the same narrowly defined product and local labor markets before the shock, but whose owners differ in their exposure to spousal unemployment risk. Balancing tests confirm that treated and control entrepreneurs are similar across a wide range of unmatched observables, including the entrepreneur's earnings, assets, revenue, leverage, and firm age.

We also implement a placebo test based on entrepreneurs whose spouses had been employed in the oil sector but *luckily* exited in the months immediately prior to the the shock. This group serves as a natural benchmark, capturing entrepreneurs whose spouses had similar prior exposure to and selection into the oil industry, but who were not subsequently exposed to the increase in labor income risk. As such, if spouses' selection into the oil industry endogenously shapes entrepreneurial outcomes, this group should capture these effects.

We use the Canadian Employer–Employee Dynamics Database (CEEDD), a linked administrative dataset covering the universe of individuals, families, and incorporated firms in Canada. A key advantage of the CEEDD is the presence of family links—often missing in other datasets—which allow us to identify entrepreneurial households and link them to detailed firm-level information and employment outcomes. For entrepreneurial firms, we observe detailed balance-sheet and income-statement information, as well as labor market outcomes for their employees. The Canadian setting is also particularly well suited to our question. The

vast majority of entrepreneurs are part of households: by age 35, more than 80 percent of entrepreneurs are married or cohabiting.[2]

We start by showing that Canadian workers exposed to the 2014 oil shock were more likely to become unemployed and more likely to experience an long-term income shock. Specifically, workers in the oil industry were 4–5 percentage points more likely to lose their jobs and approximately 2 percentage points more likely to remain unemployed. Workers in the oil industry also experienced an earnings decline of over 5 percent. These employment effects are not surprising given the magnitude of the decline in oil prices in 2014 and the well-established impact on aggregate unemployment documented in the literature. Additionally, an event plot analysis shows no pre-trends between treated and non-treated spouses, further strengthening the validity of our setting.

We then examine the firm decisions of entrepreneurs married to these workers. We first document that these entrepreneurs are more likely to reduce firm size: they are more likely to cut asset growth by over 4 percentage points, a decline of about 8 percent of the standard deviation. This reduction is also accompanied by a reduction in risk exposure: leverage declines by almost 3 percentage points. When restricting the sample to spouses who remained employed, the point estimates are almost identical. Additionally, *lucky* workers who moved from the oil to non-oil sector months before the shock are no more likely to face unemployment, and their entrepreneur spouses' firm outcomes are indistinguishable from the control group. These results show the importance of household income risk for entrepreneurs, and are consistent with precautionary saving motives as well as tighter borrowing constraints.

However, these two mechanisms have different theoretical implications for dividend disbursements. Tighter borrowing constraints predict a reduction in dividend payouts, whereas precautionary motives imply an increase in dividend disbursements. Since our data allow us to measure dividend payouts even for small firms, we use these data to disentangle the relative importance of the two mechanisms. We find that entrepreneurs with spouses in the oil industry

[2]Figure 4 provides a first piece of descriptive evidence motivating our focus on household risk. The figure plots the share of individuals who are married or cohabiting at each age for entrepreneurs and for the overall population in Canada. This figure also shows another interesting pattern. Entrepreneurs are consistently more likely to be partnered than the population at large, with an approximate 10 percentage point gap that persists throughout most of the working life.

are 3 percentage points more likely to increase dividend payouts. This effect is economically large and offers support for the role precautionary motives in explaining our findings.

We further investigate the role of risk aversion using rich demographic data in our dataset. We show that older entrepreneurs are more likely to reduce asset growth by an additional 6 percentage points relative to younger entrepreneurs. In addition, older female entrepreneurs reduce asset growth by an additional 8 percentage points. If these older female entrepreneurs also have high intra-household bargaining power, asset growth is likely to decline by 13 percentage points. As older individuals and women tend to be more risk averse, these findings are consistent with the precautionary motives mechanism.

Precautionary behavior should also be stronger when firm cash flows are riskier. More volatile firm cash flows leave entrepreneurs more exposed to downside risk, leading them to cut back investment more aggressively following an increase in household income risk. We provide evidence consistent with this prediction. Firms in low-risk industries exhibit small and statistically insignificant effects on asset growth, whereas firms in more volatile industries experience a large decline in investment following an increase in spousal income risk. Additionally, we show that the negative effect on asset growth is substantially larger among young firms.[3]

We conduct several additional robustness tests to rule alternative interpretations of our results. First, we rule out the possibility that geographical proximity drives our results. Not only do all our regressions include community zone by year fixed effects, but we also conduct an additional analysis using postal code by year fixed effects. Our results are even larger under this latter specification.[4]

Second, we also rule out that business relations with the oil industry migh explain our findings. Managers in the oil sector are no less exposed to the oil shock than other employees, and entrepreneurs married to managers experience slightly smaller, statistically insignificant declines in asset growth following the shock. If business opportunities were transmitted through

[3]We define young firms as those that are at most five years old in 2014, when the shock occurs.

[4]These larger results may be driven by the sample of firms used for identification. The inclusion of these strict fixed effects reduces the sample of comparable firms, since postal codes cover an extremely small geographical area.

the spouse's network, entrepreneurs married to managers would experience larger post-shock declines in firm growth. Additionally, we recall that *lucky* spouses are no more likely to experience unemployment, and the firm outcomes of their entrepreneur spouses are statistically and economically indistinguishable from those of the control group. This pattern is difficult to reconcile with an explanation based on persistent business relationships with the oil sector: if our estimates were driven by such relationships, it is unlikely that exposure would suddenly disappear after a spouse's industry switch.

Third, we also assess whether informational spillovers within the household can account for our results. We first note two patterns that are difficult to reconcile with a pure information story. First, the estimated effects persist through 2016, well after the 2014 oil price collapse had become common knowledge in Canada. Second, if informational advantages were central, the effects should be weaker in oil-producing provinces—where oil-sector developments are more salient and widely monitored—yet we find the opposite, with larger effects in oil provinces. Additionally, we split the control group into firms that employ a top non-owner manager (i.e., one of the highest-ranked employees) whose spouse works in the oil industry and firms that do not. If the main estimates were driven by information, then providing the control group with a comparable informational channel should attenuate the differences between treated and control entrepreneurs. We find little evidence of such attenuation.

We develop a heterogeneous-agent general equilibrium model that captures the dual role of entrepreneurs as both firm operators and household consumers. The model compares two economies: a baseline economy (E1) in which entrepreneurs and workers are distinct agents who interact only through factor and credit markets, and a household economy (E2) in which a single agent jointly manages an entrepreneurial firm and supplies labor. Both economies feature Epstein-Zin recursive preferences, a corporate and entrepreneurial sector, and collateral-based borrowing constraints in which entrepreneurs pledge a fraction of their capital as collateral. In the household economy, agents can additionally pledge a fraction of future spousal wage income, expanding their borrowing capacity beyond what capital alone would support. Two types of entrepreneurs with different risk aversion coexist: a low type, whose portfolio tilts toward entrepreneurial income, and a high type, who relies more heavily on labor income.

We also consider idiosyncratic unemployment risk: in each period a worker faces a stochastic probability of job loss that rises during downturns.

We find that household structure reshapes the cross-sectional allocation of capital. Following Moll (2014), we show that E2 delivers greater capital concentration among high-productivity entrepreneurs than E1, as spousal wage income helps productive types scale up through borrowing and self-financing. Quantitatively, the household economy amplifies a 5% aggregate TFP shock by 1.02$\times$ relative to the baseline, entirely through a precautionary savings motive. When productivity falls and unemployment rises, E2 entrepreneurs face greater wealth risk from the joint lottery over productivity and employment, triggering deleveraging that is largely absent in E1. This response is strongest among high-risk-aversion entrepreneurs, who cut capital and borrowing much more aggressively than low-risk-aversion types, replicating our empirical findings. After the shock, wealth concentrates further among high-productivity types, improving allocative efficiency by the Moll TFP measure, but this improvement reflects distress rather than productive accumulation.

Related literature. Our work contributes to four strands of literature. The first examines how economic shocks propagate across firms and regions. A growing body of research shows that aggregate shocks spread through interconnected production and financial networks, including supply-chain relationships (Acemoglu et al., 2012; Carvalho and Tahbaz-Salehi, 2019; Carvalho et al., 2021; Barrot and Sauvagnat, 2016), multinational or multi-establishment firms (Cravino and Levchenko, 2017; Giroud and Mueller, 2019; Bena, Dinc, and Erel, 2022; Biermann and Huber, 2024), and credit linkages (Jiménez et al., 2012; Chodorow-Reich, 2014; Amiti and Weinstein, 2018). These studies show that linkages among firms and financial intermediaries amplify local shocks, scaling them up to the aggregate level. We contribute to this literature by documenting a household balance-sheet transmission in entrepreneurial families: when one spouse faces a sectoral downturn, the resulting increase in household income risk induces the entrepreneur's firm to retrench. This household-to-firm channel operates outside formal economic relationships among firms, adding a previously unmeasured propagation node through which macroeconomic shocks can spread.

The second strand of literature studies how personal balance sheets affect entrepreneurial activities. Rising house prices that increase collateral values increase business formation and investment, consistent with housing wealth relaxing borrowing constraints for entrepreneurs (Corradin and Popov, 2015; Adelino, Schoar, and Severino, 2015; Schmalz, Sraer, and Thesmar, 2017). Similarly, cash windfalls, such as inheritances and lottery winnings, lead to higher firm entry rates, (Hurst and Lusardi, 2004; Huang et al., 2023; Bermejo et al., 2025). Existing studies link firm outcomes to entrepreneurs' personal liquidity, but typically treat the entrepreneur in isolation. We extend this literature by recognizing that most entrepreneurs are part of households whose combined balance sheets determine financial capacity. The spouse's exposure to labor-market fluctuations thus represents an important and largely overlooked source of household income risk with direct implications for firm behavior.

Third, our paper contributes to a large literature that examines how households respond to income and unemployment risk by building precautionary savings and rebalancing their financial portfolios toward safer or more liquid assets (Carroll, Dynan, and Krane, 2003; Bayer et al., 2019; Catherine, 2022; Juelsrud and Wold, 2025). This literature typically focuses on households that hold diversified financial assets. In contrast, many entrepreneurial households invest most of their resources in their own private businesses, resulting in highly undiversified portfolios that expose household wealth to firm performance (Moskowitz and Vissing-Jørgensen, 2002). We fill this gap by studying how such households adjust firm policies in response to changes in income risk.

Finally, our paper relates to the literature on risk sharing within households. Under incomplete markets, couples smooth consumption by adjusting labor supply and savings when facing negative income shocks (Hess, 2004; Blundell, Pistaferri, and Saporta-Eksten, 2016; Mankart and Oikonomou, 2017). One well-documented adjustment margin is the *added-worker* effect, whereby a married woman enters the labor force in response to her husband's job loss (Stephens, 2002; Guner, Kulikova, and Valladares-Esteban, 2025). We show that intra-household insurance can have aggregate implications in entrepreneurial families because when the spouse of a business owner faces higher unemployment risk, the entrepreneur responds by altering firm policies. These corporate adjustments transmit household risk-management

motives beyond the family, affecting firms and workers and amplifying the propagation of macroeconomic shocks.

2 Institutional Setting

2.1 The 2014 oil price shock in Canada

The oil and gas sector is a cornerstone of the Canadian economy, contributing 7.5% to Canada's GDP in 2013 (Fremeth, Holburn, and Loudermilk, 2016). Canada's crude oil production is heavily concentrated in a few provinces, with Alberta alone accounting for more than 75% of oil production, followed by Saskatchewan, Newfoundland and Labrador, and British Columbia. We will refer to these provinces as the oil-exposed regions. Other provinces are also economically linked to the oil sector through upstream supply through upstream activities such as construction and specialized industrial services. In mid-2014, global oil prices collapsed abruptly, falling from over USD 100 per barrel in June to below USD 50 by early 2015. The decline triggered a sharp contraction in the regions most exposed to oil activity. Employment in oil and gas extraction fell by nearly 20% in the affected provinces.

Figure 5 summarizes how regional macroeconomic conditions evolved around this episode. Panel A plots real GDP per capita and shows that oil-exposed provinces expanded rapidly before the shock. Between 2010 and 2014, GDP per capita in these provinces rose by about 11%, compared with 7% nationally. This reflects a period of high investment in the oil sands and in offshore production, alongside strong demand for related industrial support. After mid-2014, the pattern reverses sharply. When WTI prices fall below 50 USD, GDP per capita in oil-exposed regions declines by roughly 6% between the end of 2014 and mid-2016. The national decline over the same period is only 1.5%. Economic activity in these provinces moves closely with the oil price and displays substantially greater volatility than the national series.

Panel B shows the corresponding evolution of unemployment. Prior to the shock, unemployment in oil-exposed provinces remains below 6%, while the national rate is slightly above 7%. These low regional rates are consistent with strong demand for workers in extraction and in the construction and industrial services that accompany large-scale resource development. Once

oil prices collapse, unemployment in the oil-producing provinces rises rapidly and eventually exceeds the national rate. The timing of this increase closely matches the fall in oil prices.

Figure 6 presents the geographic distribution of these labor-market patterns across Canadian Economic Regions. Panel A displays unemployment rates in 2013, before the shock. Unemployment is lowest in regions directly connected to oil production or to the industrial and construction activities that support it. Alberta's economic regions, including those containing the oil sands and the major project corridors, are uniformly below 5%. Saskatchewan and the resource-intensive regions of northern British Columbia show a similar pattern. In Newfoundland and Labrador, unemployment is higher overall, but the economic region most closely tied to offshore oil activity still records unemployment below 8% on the eve of the price decline.

Panel B shows unemployment rates in 2016, when oil prices reach their trough. Unemployment rises sharply in the regions most exposed to the industry. Alberta experiences the largest increases, with both the northern oil sands region and Calgary climbing above 9%. Saskatchewan and northern and western British Columbia experience notable increases as well. Newfoundland and Labrador also sees unemployment rise in its offshore areas, despite beginning from a higher baseline. Taken together, the shift from Panel A to Panel B shows that the 2014 price crash translated into pronounced and geographically concentrated labor market distress in the regions most dependent on oil activity.

2.2 Labor market institutions and the cost of job loss

The contraction that followed the oil price collapse rapidly translated into widespread layoffs in part because Canada has a flexible labor market. Employers face few restrictions on dismissals: termination without cause is permitted so long as minimum notice or pay in lieu is provided (typically two to eight weeks). There is no statutory severance regime or just-cause requirement comparable to those in many European countries. This institutional environment makes layoffs relatively inexpensive and allows firms to adjust employment quickly in response to adverse shocks. [5]

Canada's unemployment insurance system provides partial income support following sep-

[5] See OECD (2019) for a detailed review of Canadian employment protection relative to other OECD countries.

aration. Eligibility requires recent contributions to the unemployment insurance system and sufficient insurable hours. The average net replacement rate in the first months of unemployment is approximately 55-60%, slightly above the OECD average, but benefit durations tend to be shorter (OECD, 2015). As a result, displaced workers face substantial income risk, especially when reemployment prospects are weak – as in the oil-exposed provinces during the oil crash.

Despite Canada's social safety net, job loss entails large and persistent earnings losses for affected workers. Workers displaced from stable employment experience average earnings declines of about 20–30% five years after separation; losses are particularly large among older and high-seniority employees (Morissette, Zhang, and Frenette, 2007; Bonikowska and Morissette, 2012). The *scarring effect* of unemployment is well-documented and stems from a combination of factors, including the destruction of firm- and occupation-specific human capital and weaker job matches in subsequent employment (Lachowska, Mas, and Woodbury, 2020; Raposo, Portugal, and Carneiro, 2021).

3 Data

We use data from the Canadian Employer–Employee Dynamics Database (CEEDD), a comprehensive administrative dataset maintained by Statistics Canada. The CEEDD links tax and employment records for the universe of individuals, families, and firms in Canada over time. It is constructed by integrating information from multiple administrative sources, including the T1 personal tax files, the T4 Statement of Remuneration, the T1 Family File, and the T2 corporate tax files. The longitudinal structure of the CEEDD allows for tracking individuals across jobs and firms across years, while also linking family members.

The T1 personal tax file contains individual-level demographic and income information, including birth year, sex, marital status, province and city of residence, and total taxable income. The T4 file provides job-level information on annual earnings paid by each employer to each worker. This allows for complete reconstruction of employment histories and income sources. In addition, the T4 Record of Employment, which employers are required to file whenever a worker experiences an interruption in earnings, contains the reason for separation. This is particularly important in our context, as it allows us to identify spouses who were laid off due

to a shortage of work, rather than voluntarily leaving their job. The T1 Family File (T1FF) identifies household relationships by linking individuals to their married or cohabiting spouse and dependent children based on tax co-filing and benefit records. We use these links to identify each entrepreneur's spouse and connect them to their own job-level data, including earnings, employer, and industry of employment.

The T2 corporate tax file contains balance sheet and income statement data for all incorporated firms in Canada, including total assets, sales, payroll, and detailed industry codes. These files are linked to T2 Schedule 50, which reports the identity and ownership share of all shareholders holding at least 10% of the firm's equity in each year for Canadian-controlled private corporations. This feature of the data allows us to identify entrepreneurs directly as equity owners of the firm, while in the literature entrepreneurs are often proxied by the top earners in a firm. Such proxies are prone to misclassification: high earners may simply be early employees or managers without an equity stake, while true owners may compensate themselves through dividends rather than salary, making them invisible in payroll-based measures. By observing actual equity stakes, we can unambiguously distinguish entrepreneurs from employees.

4 Empirical Methodology

Measuring how firms respond to household income risk requires a source of variation that is external to the firm and affects only the entrepreneur's household—without influencing firm fundamentals. This is challenging, because in most entrepreneurial households, income risk is tightly linked to the firm itself. To isolate shocks that are not confounded by firm performance, we focus on entrepreneurs whose spouses are employed outside the firm and are not involved in its operations. This allows us to study how business decisions respond to changes in the household's broader financial environment.

A natural approach might be to use shocks to the spouse's firm or industry—such as bankruptcies, mass layoffs, or sectoral downturns. However, these shocks are often gradual and at least partly anticipated by workers, especially those employed in the affected sectors. Anticipation can influence precautionary behavior even before the shock is observed, making it harder to isolate the impact of uncertainty itself. In addition, firm- or industry-level shocks may

reflect local economic trends or be correlated with unobserved characteristics of the spouse's employment situation.

Using realized layoffs is also problematic: job loss is a consequence of underlying risk, not a measure of risk itself. It reflects the resolution of uncertainty, rather than changes in perceived or anticipated volatility. Moreover, layoffs are often non-random, concentrated among workers with lower performance or tenure, and may trigger other household responses unrelated to income risk per se.

To address these limitations, we exploit the 2014 collapse in global oil prices, which generated a sudden and unexpected rise in income risk for households with ties to the oil sector. For our purposes, it is crucial that the 2014 oil-price plunge was exogenous and unanticipated from the vantage of Canadian firms. The collapse was driven by global forces, particularly a surge in U.S. shale oil output and waning demand in emerging markets, and was exacerbated by a surprise decision in November 2014 by the Organization of the Petroleum Exporting Countries (OPEC) to maintain production despite falling prices (Arezki and Blanchard, 2015). This strategic move deviated from past policy norms and took markets off guard. There is consensus that the price collapse was largely unexpected by most forecasters (Kilian and Baumeister, 2016). As late as in October 2014, after the price decline had already started, the International Monetary Fund predicted that the oil price would stabilize: *"the average price of oil will be $102.76 a barrel in 2014 and $99.36 a barrel in 2015 and will remain unchanged in real terms over the medium term"* (IMF, 2015). This combination of exogeneity and surprise makes the oil price collapse a credible source of variation in household income risk for our identification purposes.

4.1 Construction of treatment and control groups

We start by restricting the sample to incorporated firms with at least $5,000 in total assets in 2013 to ensure that they are operating businesses rather than dormant entities or holding companies. In addition, we require that both the entrepreneur and the spouse contribute materially to household income. Specifically, we exclude households in which the entrepreneur accounts for less than 5% or more than 95% of total household income in 2013. This restriction ensures

that the household is meaningfully exposed to income risk from both spouses: if the working spouse contributes almost nothing to household earnings, their job loss is unlikely to affect entrepreneurial decisions, while if the entrepreneur contributes almost nothing, there is little scope for firm-level adjustments to buffer household shocks.

In addition, we impose several additional restrictions to ensure that the household's exposure to the oil price shock is exogenous to the entrepreneur's own business activity. We exclude any households in which one or more of the following conditions occur: (i) either the entrepreneur or the spouse own equity in the spouse's employer firm; (ii) the spouse owns equity in or receives compensation from the entrepreneur's firm; (iii) the entrepreneur receives compensation from the spouse's firm, or any other firm in the oil sector. These restrictions eliminate cases in which the entrepreneurial firm is directly exposed to oil sector shocks through the entrepreneur's own employment or spousal business integration, allowing us to isolate the impact of spousal employment in the oil sector on household income risk.

Our identification strategy compares entrepreneurs whose spouses were exposed to an exogenous increase in unemployment risk following the oil price collapse to observationally similar entrepreneurs whose spouses were not exposed. The treated group consists of entrepreneurial firms whose owner's spouse was employed in the oil sector in both 2013 and 2014. This requirement ensures that spouses entered the sector prior to the shock and were plausibly exposed to its consequences at the onset of 2014. We define the oil sector to include oil and gas extraction, support activities for oil and gas operations, and closely related upstream supplier industries. We also require that couples are observed in both 2013 and 2014, so that the entrepreneur is linked to an exposed spouse at the time of the shock. Subsequent job mobility or changes in household composition are treated as outcomes of the shock rather than selection criteria.

The control group consists of entrepreneurial firms whose owner's spouse was employed in a non-oil sector in both 2013 and 2014 and who remained married to the entrepreneur in both years. Thus, treated and control spouses have similar pre-shock employment stability, sectoral attachment, and household structure, differing only in exposure to the oil price collapse.

We implement one-to-one matching to construct a comparable control group. Matching is

performed in 2013, prior to the oil price collapse, using a caliper matching algorithm (Stepner and Garland, 2017). Caliper matching constructs comparable treatment and control groups by restricting matches to pairs whose differences in selected pre-shock characteristics fall within a pre-specified tolerance (the "caliper"). We adopt a deliberately parsimonious matching strategy based on four pre-shock characteristics: we match exactly on 4-digit industry, entrepreneur sex, and commuting zone,[6] and with a 25-log-points caliper on spousal earnings. The commuting zone match is particularly important given the geographic concentration of the oil sector. Most oil-producing activity is located in a small number of regions, and local entrepreneurial firms may be indirectly affected by the shock through general equilibrium channels such as changes in local income, labor market conditions, or demand for goods and services.

4.2 Balancing tests

Table 3 reports pre-shock characteristics for treated and control firms in 2013. Columns 1-2 show that after matching, treated and control firms are highly comparable along a broad set of household, firm, and entrepreneur characteristics, including several dimensions that are not explicitly targeted by the matching procedure. Entrepreneurs and their spouses in treated households earn an average of $99,000 and $107,000 respectively, compared with $102,000 and $104,000 among controls. This near parity shows that both partners contribute approximately equally to household income. Firm-level characteristics are equally balanced: average revenues are roughly $640,000 to $650,000 and total assets are about $650,000 in both groups. Treated firms have an average net income of around $90,000, a leverage ratio of 27%, and an average age of 10 years. Entrepreneurs retain ownership stakes of 71% on average, suggesting strong personal financial exposure to firm performance. Figure 7 displays standardized mean differences between treated and controls for all continuous covariates. After matching, observable characteristics are well balanced: all standardized differences fall below conventional thresholds.

[6]Statistics Canada defines geographic units using commuting-flow patterns, grouping adjacent municipalities into integrated local labor markets. Although they are labeled as Census Metropolitan Areas (CMAs) in the data, their construction corresponds to the commuting-zone concept widely used in labor economics. We therefore refer to them as commuting zones.

For comparison, the unmatched data in columns 3-4 illustrate why matching is essential in this setting. Prior to matching, households exposed to the oil sector differ sharply from the broader population of entrepreneurs. Spouses in oil-sector households earn nearly twice as much as those in non-oil households, reflecting the substantial wage premium in oil-related jobs. These households are also far more geographically concentrated: approximately 60% of treated households are located in Alberta, consistent with the distribution of oil-sector jobs in Canada, while only about 11% of all potential control firms are located in this province. The composition of entrepreneurs also differs sharply across groups: while only about 17% of entrepreneurs in the full sample are female, the share rises to roughly 39% among oil-exposed households, because oil-sector employees are predominantly male and their entrepreneurial partners are therefore more often women.

4.3 Empirical specifications

Oil shock and unemployment risk—. To validate our identification strategy, we first show that spouses employed in the oil sector in 2014 experience a differential increase in unemployment risk due to the oil prices collapse, relative to their matched counterparts outside the oil sector. We trace spouses' labor market outcomes around the shock by estimating the following event–study specification:

$$y_{it} = \alpha_i + \lambda_c t + \sum_{\tau \neq 2013} \theta_\tau \, \mathbb{I}\{t = \tau\} \times \mathrm{Oil}_i + X_{it}'\delta + \varepsilon_{it}, \tag{1}$$

where y_{it} represents an outcome (e.g., being unemployed) for spouse i in year t. α_i and $\lambda_c t$ denote individual and commuting zone $\times$ year fixed effects, respectively. $\mathbb{I}\{t = \tau\}$ are event-study indicators and Oil_i is an indicator for whether the spouse works in the oil sector in 2014. The set of control variables X_{it} includes a polynomial for age. The error term, ε_{it}, is clustered at the individual level.

We examine two complementary indicators of job displacement. First, we measure whether a worker is laid off, defined as an employer-initiated separation due to a "shortage of work." This information comes from the Record of Employment, a form that Canadian employers are required

to file whenever an employee separates, and which records the reason for separation. Because the ROE captures layoffs directly as reported by employers, it provides a precise measure of involuntary job loss. Second, we measure whether an individual receives Employment Insurance benefits in a given year, This measure is constructed from tax return data and serves as a proxy for unemployment spells.

Transmission to entrepreneurial firms—. We next examine whether entrepreneurs adjust firm decisions in response to a sudden increase in household income risk. The analysis contrasts entrepreneurial firms whose owners are observationally similar but differ in their exposure to the 2014 oil price collapse through their spouse's employment. In treated households, the entrepreneur's spouse was working in the oil sector in 2014 and therefore directly exposed to the sharp rise in unemployment risk that followed the oil price collapse. In control households, the spouse was not employed in the sector and remained largely unaffected. This variation provides a quasi-experimental setting to identify how income risk at the household level affects firm behavior. We test for firm-level responses by estimating the following equation:

$$y_{jt} = \alpha_j + \lambda_{cmt} + \sum_{\tau \neq 2013} \theta_\tau \, \mathbb{I}\{t = \tau\} \times \text{Oil}_{i(j)} + X'_{jt}\delta + \varepsilon_{jt}, \tag{2}$$

where y_{jt} is an outcome for firm j. α_j and λ_{mct} indicate firm and industry $\times$ commuting zone $\times$ year fixed effects, and $\text{Oil}_{i(j)}$ is an indicator equal to 1 if the entrepreneur's spouse i works in the oil sector in 2014. The set of control variables X'_{jt} includes polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of shareholders. The inclusion of firm fixed effects α_j ensures that identification comes from within-firm changes over time, absorbing all time-invariant firm characteristics. The industry-by-commuting zone-by-year fixed effects λ_{qmt} flexibly account for local economic shocks and demand conditions that vary over time across narrowly defined industries and geographic areas. This is important because some sectors may have been more exposed to the oil shock in oil-producing regions. By absorbing this variation, we ensure that identification does not rely on differential exposure across sectors or cities, but rather on differential exposure across households within the same

industry and location.

Propagation to employees—. Finally, we examine whether the increase in household income risk faced by treated entrepreneurs is reflected in the labor market outcomes of the workers they employ. The sample consists of incumbent employees who were continuously employed in 2013 and 2014 by entrepreneurs in either the treated or the control group. This design allows us to observe how the shock to the entrepreneur's household affects workers who were already attached to the firm before the oil price collapse. We estimate the following event–study specification:

$$y_{lt} = \alpha_l + \lambda_t + \sum_{\tau \neq 2013} \theta_\tau \, \mathbb{I}\{t = \tau\} \times \text{Oil}_{i(l)} + X_{lt}'\delta + \varepsilon_{lt}, \tag{3}$$

where y_{lt} is an unemployment indicator for worker l in year t. The indicator $\text{Oil}_{i(l)}$ equals one if worker l is employed by an entrepreneur whose spouse i worked in the oil sector in 2014. Individual fixed effects α_l absorb all time-invariant worker characteristics. The control variables X_{lt} include a polynomial in age. Standard errors are clustered at the individual level. The coefficients θ_τ capture whether workers employed by entrepreneurs in treated households become more likely to experience unemployment relative to workers employed by entrepreneurs in the control group. The comparison is across employees who worked in similar firms, in the same local labor markets, before the shock but who differ only in whether their employer's household was directly exposed to the oil price collapse.

5 Results

5.1 Exposure to unemployment risk

We begin by documenting how the collapse in oil prices affected the labor market outcomes of spouses employed in the oil sector. Treated spouses are those working in the oil sector in both 2013 and 2014. The control group consists of matched spouses working outside the sector but residing in the same commuting zone in 2013.

Figure 8 plots event–study coefficients for layoffs and unemployment estimated from equation 1. Panels (a) and (b) show the probability of experiencing a layoff, defined as an employer-

initiated separation recorded as "shortage of work." Before 2014, the coefficients are close to zero and display no differential trend between the treated and the control groups. The pattern changes starting in 2014. The probability of being laid off increases by about 2.3 percentage points in that year and continues to rise in the following two years. By 2016, oil-sector spouses are roughly 5 percentage points more likely to experience an involuntary separation than their matched counterparts.

To capture realized unemployment, we complement this measure with information on unemployment insurance benefits from administrative tax data. These two indicators reflect different phases of job displacement. Layoffs measure employer-reported involuntary separations. Unemployment insurance receipt measures periods spent without work. Not all separations result in unemployment, and some workers transition directly into new jobs. Using both indicators provides a more complete picture of exposure to job loss. Panels (c) and (d) show the effects on unemployment insurance receipt. As with layoffs, the pre-shock coefficients are close to zero. The effect in 2014 is small, consistent with the aggregate evidence that unemployment rates in oil-producing provinces remained very low through most of that year. The probability of receiving unemployment benefits rises only in 2015 and 2016, reaching roughly 2.5 percentage points by the end of the window. This pattern reflects the fact that early separations could often be absorbed by the local labor market, with unemployment rising only as conditions deteriorated further.

Table 4 presents the average treatment effects on the treated comparing the period before the shock (2010 to 2013) to the shock period (2014 to 2016). Each panel reports estimates from increasingly saturated geographic fixed–effects specifications. Columns (1) to (3) use the full national sample, and columns (4) to (6) restrict the sample to the four oil-exposed provinces. All columns include individual fixed effects and control for a polynomial in age. Column (1) includes year fixed effects; column (2) replaces year fixed effects with province-by-year effects, and column (3) with commuting-zone-by-year effects.

Panel A shows that the probability of experiencing an involuntary separation increases by about 4 percentage points in the full sample across all three specifications. In the oil-region subsample, the estimates are larger at about 5.5 percentage points. These patterns are

highly stable across specifications and consistent with the concentration of layoffs in the most exposed provinces. Panel B shows that treated spouses are about 2 percentage points more likely to receive unemployment insurance benefits. The estimates are smaller than for layoffs because the post period includes 2014, a year in which the event–study coefficients remain near zero. This mechanical averaging reduces the magnitude of the estimated average treatment effect. Panel C reports the effects on log annual earnings. In the full sample, treated spouses experience earnings losses of about 4.6 to 4.9 percent relative to their matched counterparts. In the oil-region subsample, the losses are larger at roughly 5.5 percent. As with the other outcomes, the estimates vary little across specifications.

These results indicate that spouses employed in the oil sector experienced a substantial increase in job loss and unemployment risk and a decline in earnings following the 2014 oil price crash. These patterns indicate that treated households experience an economically important increase in income risk, setting the stage for examining how entrepreneurs adjust their firms' decisions in response.

5.2 Do entrepreneurs internalize household shocks?

The evidence above establishes that spouses employed in the oil sector faced a sudden and unexpected increase in the risk of job loss after the 2014 price collapse. We now turn to the central question of the paper: do entrepreneurs internalize this household shock when making business decisions? If entrepreneurs treat their firms as entirely separate from the household, spousal unemployment risk should not affect firm policies. But if household resources are pooled, then greater uncertainty in the spouse's labor market position may prompt the entrepreneur to adjust, investment, payout, or capital structure decisions in order to buffer overall household income risk.

Figure 9 examines how firm policies adjust in response to the household shock using equation 2. All specifications include firm fixed effects and industry × commuting-zone × year fixed effects, so identification comes from within-firm changes relative to other firms in the same industry and local labor market. For each outcome, the first panel reports estimates without additional controls, and the second panel adds controls for firm age, entrepreneur

age, marital status, and the number of owners. Panels (a) shows that asset growth evolves in parallel for treated and control firms until 2013. In 2014, asset growth for treated firms falls by approximately 4 percentage points relative to controls and remains roughly 5 percentage points lower by 2016. The estimates change very little when adding the control variables in panel (b). This slowdown is consistent with the prediction that households respond to higher income risk by reducing firm scale.

Panels (c) shows that dividend payout also responds immediately. Before 2014, the probability of distributing dividends is similar across the treated and control groups. In 2014, treated firms become approximately 3 percentage points more likely to distribute cash, and this difference remains similar at 3.5 percentage points through 2016. In panel (d), which includes controls, the difference in 2016 is closer to 3 percentage points. The fact that payouts rise at the same time that expansion slows suggests a reallocation of resources away from investment and toward liquidity extraction, consistent with precautionary motives at the household level.

Panels (e) and (f) report the effects on financial leverage. In 2014, treated firms reduce leverage by about 2 percentage points relative to controls; the gap widens to roughly 3 percentage points in 2015 before narrowing again in 2016, when the difference is closer to 2 percentage points and less precisely estimated. The direction of the effect is informative because leverage could, in principle, move either way. On the one hand, entrepreneurs facing greater household income risk may prefer to rely more on external financing and commit less of their own equity to the firm. On the other hand, higher risk on the household side can tighten borrowing limits if lenders perceive the entrepreneur as less able to repay. The decline in leverage suggests that financing constraints dominate.

Table 5 presents the corresponding average treatment effects (ATT). Column (1) presents the baseline specification with firm fixed effects and industry–year fixed effects, comparing treated and control firms within narrowly defined sectors. Column (2) adds control variables. Columns (3)–(4) adds commuting-zone–year effects. Columns (5)–(6), the most demanding specifications, include the full set of industry × commuting-zone × year fixed effects, absorbing any time-varying shocks at the intersection of sector and local economy. Columns (7)–(8) repeat these two specifications for firms in the four oil-producing provinces.

In Panel A, the estimated treatment effect on asset growth ranges from –0.034 to –0.043 in the full sample and from –0.054 to –0.059 in the oil-province subsample. Given a standard deviation of 0.51 in 2013, the implied effects correspond to 7–12 percent of a standard deviation. Panel B shows the effects on dividend disbursement. The estimated increase in payout ranges from 0.024 to 0.030 percentage points in the full sample and from 0.036 to 0.040 in the oil provinces. These effects correspond to a 5–8 percent increase in the likelihood of distributing dividends relative to the mean. Panel C presents the effects on leverage. Treated firms reduce leverage by roughly 2.5–3 percentage points, with slightly larger declines in the oil-province subsample. Benchmarking against the 2013 standard deviation of 0.52 implies a decline of 5 percent of a standard-deviation.

Table 6 examines how firm operations respond to the shock. We first consider revenues and expenses to assess whether treated firms experience an immediate contraction in activity. Columns 1-4 show that revenue growth declines by x percentage points, while expenses show a small and statistically insignificant response. We then examine whether firms become less productive. In columns 5-6, we find that profit margins decline by x percentage points. Firm productivity, measured as revenues per employee, falls by x percentage points (columns 7-8). These effects are estimated less precisely than the effects we documented earlier on investment and dividend disbursement. Part of the reason is their timing: we find little evidence of changes in firm operations in 2014; declines in revenue growth, profitability, and productivity only emerge in 2015 and 2016. The absence of immediate effects suggests that firms were not directly affected by a sudden deterioration in business conditions. Instead, the results suggest that entrepreneurs gradually constrain firm growth, leading to slower growth and lower operating efficiency.

Finally, we examine firm survival. Survival is commonly interpreted as a proxy for risk-taking: entrepreneurs who pursue more risky projects face a higher likelihood of failure, while more risk-averse owners tend to operate firms that survive longer. Entrepreneurs exposed to greater household income risk may become more cautious and reduce risk-taking, which would predict an increase in the likelihood of firm survival. On the other hand, they also have the outside option of shutting down the business and returning to wage employment, which may be

perceived as a safer and more predictable source of income. In that case, firm exit may reflect a deliberate reallocation away from entrepreneurial risk rather than failure. Consistent with this latter channel, we find a small but statistically significant decline in firm survival among treated firms (columns 9-10).

Unemployment risk vs. realized job loss—. The distinction between increased income risk and realized income shocks is central to interpreting our findings. If entrepreneurs adjust firm policies only after a spouse loses their job, the observed responses could reflect the direct consequences of lower household income. By contrast, adjustments that occur before any job loss indicate a response to heightened risk rather than to realized shocks. A key advantage of our setting is that it allows us to examine whether entrepreneurs respond even in the absence of realized job loss.

To test this, we restrict the sample to households in which both treated and control spouses remain continuously employed through 2016. Figure A.1 shows that among these households, we still observe adjustments in firm policy starting in 2014. Firms reduce asset growth, increase dividend payouts, and lower leverage in patterns that are very similar to the baseline results. Since layoffs are not randomly assigned, this exercise, in which we condition on realized employment outcomes, is not designed to compare households that do and do not experience layoffs. Instead, it shows that our main findings are not driven by the subset of households that actually experience job loss and are consistent with precautionary responses to elevated perceived income risk.

Transmission to employees—. Figure 10 examines whether the oil-price shock affecting entrepreneurs' households also affects the workers they employ. The sample is restricted to employees who were continuously employed by the same entrepreneur in 2013 and 2014, so the estimates capture how the shock influences an ongoing employment relationship rather than changes in workforce composition. The event–study coefficients show no differences between treated and control workers before 2014. After the oil price shock, unemployment among workers in treated firms begins to rise gradually. In 2014, unemployment among workers in treated firms rises by 1.3 percentage points relative to the control group. The gap widens further

in 2015 and 2016, reaching close to 2.7 percentage points by the end of the sample. The pattern indicates a steady transmission of risk: as treated entrepreneurs face greater income uncertainty, their employees become increasingly exposed to unemployment.

Table 7 reports the corresponding average treatment effects. Columns (1)–(2) include worker and year fixed effects; columns (3)–(4) replace year effects with commuting-zone–by–year effects, comparing workers who remain in the same local labor market. Across these specifications, the estimated Post × Oil coefficient ranges from 0.014 to 0.017, indicating a stable and statistically precise increase in unemployment risk for workers employed by treated entrepreneurs.

Columns (5)–(6) examine heterogeneity by worker age. The additional increase in unemployment is concentrated among younger employees: the coefficient on the triple interaction term indicates that workers under 30 in treated firms face an extra rise of about 2 percentage points in unemployment relative to similar workers in control firms. This pattern might reflect the fact that younger employees are less costly to dismiss as notice and severance obligations grow with tenure. As a result, when entrepreneurs face higher household income risk, young workers become the natural adjustment margin. Since younger workers also tend to have steeper productivity profiles, their displacement provides a plausible channel through which household-level shocks may contribute to subsequent declines in firm productivity, consistent with Caggese, Cuñat, and Metzger (2019).

5.3 Mechanisms

Precautionary motives—. The results above show that entrepreneurs adjust firm policies following an increase in their spouse's unemployment risk. A natural interpretation is that entrepreneurs treat the firm as part of the household's balance sheet and adjust firm decisions in response to changes in household risk exposure. In this framework, spousal labor income risk constitutes a source of non-diversifiable background risk, raising the marginal value of safe and liquid assets relative to risky firm cash flows.

As household risk increases, optimal behavior involves rebalancing the household portfolio away from risky business income and toward precautionary savings. The increase in dividend

disbursements documented in the previous section is consistent with this mechanism: dividend payouts allow entrepreneurs to transfer resources out of the firm and into safer household assets in response to elevated income uncertainty. Importantly, this pattern is difficult to reconcile with an explanation based solely on tighter borrowing constraints, which would instead predict the retention of internal funds and a reduction in dividend payouts.

A key implication of precautionary behavior is that responses to background risk should vary systematically with the entrepreneur's degree of risk aversion. As we show in our theoretical framework, entrepreneurs who are more risk averse place greater weight on downside household risk and therefore reduce investment more aggressively when spousal income becomes more uncertain. Table 8 tests this prediction by examining heterogeneity in asset growth along demographic characteristics that have been shown to correlate with risk aversion.

Columns (1) and (2) examine age heterogeneity. We find that for entrepreneurs in the upper quartile of the age distribution in 2014, asset growth declines by an additional 5.8-6.2 percentage points relative to younger entrepreneurs. This result is consistent with experimental and survey-based evidence that risk tolerance declines with age, reflecting shorter horizons and lower tolerance for downside outcomes (Dohmen et al., 2011; Falk et al., 2018). Columns (3) and (4) focus on older female entrepreneurs. In this group, the contraction in asset growth is markedly stronger. Relative to the baseline effect of approximately 3–4 percentage points, asset growth declines by an additional 8-8.3 percentage points among women in the top age quartile, consistent with gender differences in risk tolerance (Charness and Gneezy, 2012). Finally, columns (5) and (6) show that the response is strongest among older female entrepreneurs with high intra-household bargaining power, proxied by having an above-median share of cumulative household income. In this group, asset growth falls by an additional 12.4 to 13.2 percentage points following the shock. This magnitude is more than three times the baseline effect, consistent with the idea that precautionary motives are amplified when the entrepreneur both bears more household risk and has greater control over household financial decisions.

In our theoretical framework, precautionary behavior is stronger when firm cash flows are riskier. Thus, responses to background risk should also depend on the riskiness of the entrepreneur's firm. When firm cash flows are volatile, increases in household income risk

make entrepreneurs more exposed to downside risk, leading them to cut back investment more aggressively. In contrast, when firms generate sufficiently stable cash flows, additional household risk need not trigger precautionary behavior. In the limit, a perfectly safe firm would serve as a hedge, implying no incentive to reduce (and potentially an incentive to increase) exposure to the firm.

Table 9 examines this prediction by studying heterogeneity in asset growth across firms operating in industries with different levels of risk. We proxy for firm risk using within-firm ROA volatility estimated using the universe of firms in the CEEDD. Industries are classified as high- or low-risk based on a median split of this industry-average volatility measure. Firms in low-risk industries (columns 1-4) exhibit small and statistically insignificant effects on asset growth, with coefficients ranging between 1 and 2.5 percentage points across specifications. By contrast, columns 5-8 show that the effects are concentrated among firms in more volatile industries, in which firms experience a decline in investment of about 5 to 7 percentage points following the increase in spousal income risk.

We also consider firm age as an additional dimension along which firm-level risk varies. Younger firms tend to have more volatile cash flows than established firms, making precautionary responses to background risk more pronounced. Table A.1 shows that the decline in asset growth following the shock is substantially larger among young firms, defined as firms that are at most 5 years old in 2014, when the shock occurs. In this group, asset growth falls by roughly two to three times as much as in more established firms: in our most demanding specification, which includes firm-level controls and industry-by-commuting zone-by-year fixed effects, investment in young firms declines by 8.9 percentage points vs. 3.2 percentage points for older firms.

Firm linkages to the oil sector—. A second potential mechanism driving our results might be that entrepreneurs with a spouse employed in the oil sector may also operate firms that are themselves directly linked to oil companies. For example, through geographic proximity or supply chains. If such firms were more exposed to revenue disruptions following the oil price collapse, the observed firm-level adjustments could, in principle, reflect direct business shocks rather than a household-level response to elevated income risk. Our baseline specification

substantially limits the scope for this mechanism by comparing treated and control firms within narrowly defined 4-digit industry × commuting zone cells. This design restricts comparisons to firms operating in the same narrowly defined product market and within the same local labor market, holding constant core aspects of business operations and local economic conditions.

In table A.2 we test the effect of geographic proximity by re-estimating equation 2 using postal codes rather than commuting zones, thereby tightening the geographic comparison to extremely small spatial units. Canadian postal codes are six-character identifiers that capture very fine locations: in urban areas, a single code often corresponds to one block face or even a single large building. In the administrative data, these identifiers are fully anonymized, so their precise coordinates, boundaries, and contiguity are not known, but they remain useful as high-resolution location markers. By including postal code $\times$ year fixed effects, we compare treated and control entrepreneurs whose firms are effectively located in the same small area, substantially reducing the possibility that differential exposure to nearby oil-related businesses drives the results. Because postal codes define such narrow geographic cells, the estimation sample falls by more than 85%. Within this much smaller sample, the estimated effects are larger in magnitude (12.6 to 14.1 percentage points) but less precisely measured. The pattern is consistent with our baseline findings and reinforces the conclusion that hyper-local business linkages to the oil sector are unlikely to account for the observed adjustments.

A second potential channel through which entrepreneurial firms could be indirectly exposed to oil-sector conditions is the flow of business opportunities through the spouse's employer. If entrepreneurs obtain contracts or customers through their spouse's workplace, this channel should be strongest when the spouse occupies a prominent position within her firm. We therefore focus on spouses who are managers, defined as workers in the top 5% of the earnings distribution within their firm, and, in particular, on managers with above-median tenure, whom we refer to as entrenched managers.

Table 10 provides no support for this mechanism. First, we show in Panel A that managers in the oil sector are similarly exposed to the oil shock as other employees; the difference in the probability of losing their job, ranging from 1 to 1.4 percentage points, is not statistically

significant.[7] Panel B then examines asset-growth responses among entrepreneurs whose spouses are managers. If business opportunities flowed through the spouse's network, entrepreneurs married to managers-especially entrenched managers—should experience larger declines in firm growth following the shock. Instead, we find the opposite pattern: entrepreneurs with manager spouses exhibit slightly smaller declines in asset growth, ranging from 0.4 p.p. for all managers to 1.4 p.p. for entrenched managers, although the differences are statistically insignificant.

To further assess the possibility that firms owned by spouses in the oil sector might have business linkages to the oil sector due to contracts or proximity, in Table A.3 we conduct a placebo test that uses the same matching procedure as in the main analysis but compares entrepreneurs whose spouses leave the oil sector in the first half of 2014 to entrepreneurs whose spouses leave other sectors at the same time. Because both groups experience a spousal job transition before the oil price shock, neither faces a differential change in household income risk (as we show in columns 1-3). Any existing business ties between the entrepreneur's firm and oil companies should persist in the short run when the spouse leaves their job; thus, if our main estimates were driven by direct business ties to the oil sector, we would expect to see similar patterns in this placebo comparison. Instead, we find no differential change in asset growth between the two groups. Taken together, these results suggest that business opportunities transmitted through the spouse's workplace are unlikely to drive the investment responses we document.

Information transmission—. Finally, a potential mechanism is that treated entrepreneurs may receive earlier or clearer signals about the impending downturn in the oil sector, consistent with information spillovers between spouses (Delavande, Koşar, and Zafar, 2025). If spouses employed in oil learn sooner that a major contraction is coming, entrepreneurs might reduce investment in anticipation of weaker aggregate conditions rather than in response to increased household income risk. Two patterns in the data are difficult to reconcile with this interpretation. First, the estimated effects persist through 2016, well after the collapse in oil prices had become widely known and incorporated into public expectations. If advance information were driving

[7] This pattern is consistent with evidence that boards do not insulate managers from aggregate shocks (Jenter and Kanaan, 2015).

the results, the effects should fade once the shock is fully understood by all firms. Second, the effects are larger in oil-producing provinces, where the public is more aware and closely monitoring the sectoral developments related to oil prices. An information-based mechanism would predict the opposite pattern.

A further implication of an information-based channel is that the effects should be strongest when the spouse is in a position to learn and transmit information about the employer and the sector. This interpretation is further weakened by the evidence on manager spouses presented in Table 10: entrepreneurs married to managers, who should be best positioned to transmit information, do not exhibit stronger investment responses. Table A.4 provides additional evidence against the information-based mechanism. If treated entrepreneurs adjust more strongly because they have access to timely signals through the spouse's job, then the gap between treated and control firms should narrow when control entrepreneurs also have access to similar information. To test this, we identify control firms that employ a top non-owner worker whose spouse works in the oil sector. These workers, who represent the firm's highest-ranked employees, are the most plausible conduit for transmitting industry-related news or expectations about economic conditions to the entrepreneur. If advance information were central, the treatment effect should be smaller when the control firms employ such employees. The results in Table A.4 provide little evidence of this attenuation, suggesting that differential access to information about the coming downturn is unlikely to account for the main findings.

6 Macro Model

We develop a heterogeneous-agent general equilibrium model that captures the dual role of entrepreneurs as both firm operators and household consumers. The model compares two economies: a baseline economy (E1) in which entrepreneurs and workers are distinct agents who interact only through factor and credit markets, and a household economy (E2) in which a single agent jointly manages an entrepreneurial firm and supplies labor. Both economies feature Epstein-Zin recursive preferences, a corporate and entrepreneurial sector, and collateral-based borrowing constraints in which entrepreneurs pledge a fraction of their capital as collateral. In the household economy, agents can additionally pledge a fraction of future spousal wage

income, expanding their borrowing capacity beyond what capital alone would support. Two types of entrepreneurs with different risk aversion coexist: a low type, whose portfolio tilts toward entrepreneurial income, and a high type, who relies more heavily on labor income. The household economy also introduces idiosyncratic unemployment risk: in each period the entrepreneur's spouse faces a stochastic probability of job loss that rises during downturns, creating a compound lottery over both productivity and employment shocks.

6.1 Environment and Preferences

The economy is populated by entrepreneurs (mass N_E) and workers (mass N_W). Entrepreneurs are divided into two types: an L-type with low risk aversion γ_L and an H-type with high risk aversion γ_H, present in equal proportions. All agents have Epstein-Zin recursive preferences, which decouple risk aversion from the intertemporal elasticity of substitution. The value function aggregator is:

$$V_t = \left[(1-\beta)\, c_t^{\rho} + \beta\, \mathrm{CE}_t^{\rho}\right]^{1/\rho}, \tag{4}$$

where $\rho = 1 - 1/\psi$, ψ is the IES, and β is the discount factor. The certainty equivalent is:

$$\mathrm{CE}_t = \left[{}_t\left(V_{t+1}^{1-\gamma}\right)\right]^{1/(1-\gamma)}, \tag{5}$$

where γ is the coefficient of relative risk aversion. Importantly, entrepreneurs maintain the *same* risk aversion γ_i ($i \in \{L, H\}$) in both E1 and E2. The household structure does not alter preferences; it only changes the budget constraint and the set of risks the entrepreneur internalizes.

6.2 Corporate Sector and Wage Determination

A competitive corporate sector operates a constant-returns-to-scale Cobb-Douglas technology:

$$Y_C = A \cdot K_C^{\alpha} \cdot N_C^{1-\alpha}, \tag{6}$$

where A is aggregate TFP ($A = 1$ in baseline, $A = 0.95$ under the shock). Corporate capital K_C is endogenous, pinned by a no-arbitrage condition requiring that the return on corporate capital equals the bond rate: $r_K = r$. This condition, combined with the corporate sector's first-order conditions, pins the capital-labor ratio and the wage directly from the interest rate:

$$\frac{K_C}{N_C} = \left(\frac{r+\delta}{\alpha A}\right)^{1/(\alpha-1)}, \tag{7}$$

$$w = (1-\alpha)\, A \left(\frac{K_C}{N_C}\right)^{\alpha}. \tag{8}$$

This unified capital market structure follows Cagetti and De Nardi (2006): a single interest rate r governs both the return on corporate capital and the cost of entrepreneurial borrowing, allowing savings to flow across both uses. Workers receive their capital income through the return r on their savings, which finance both corporate capital and entrepreneurial borrowing. Since the CRS corporate sector earns zero economic profit by Euler's theorem, no separate dividend is needed. Once the equilibrium labor allocation determines N_C, corporate capital is recovered as $K_C = (K_C/N_C) \times N_C$.

6.3 Entrepreneurial Sector

Each entrepreneur operates a decreasing-returns-to-scale technology, following the span-of-control formulation of Lucas (1978):

$$y = \theta \cdot A \cdot z \cdot \left(k^{\alpha} n^{1-\alpha}\right)^{\nu}, \tag{9}$$

where θ is a scale parameter, z is idiosyncratic productivity, and $\nu < 1$ governs decreasing returns. Labor demand is optimized out analytically:

$$n^* = \left(\frac{\nu(1-\alpha)\theta A z\, k^{\nu\alpha}}{w}\right)^{1/(1-\nu(1-\alpha))}, \tag{10}$$

yielding firm profits $\pi(k,z) = y(k,n^*,z) - w \cdot n^*$. Entrepreneurial productivity follows an AR(1) in logs: $\ln z_{t+1} = \varphi_z \ln z_t + \varepsilon_t$, with $\varepsilon_t \sim N(0, \sigma_z^2)$.

6.4 Collateral Constraints

Entrepreneurs choose how much to borrow, up to a collateral limit. In E1 (standalone):

$$b \leq \theta_k(1-\delta)\, k'. \tag{11}$$

In E2, the household can additionally pledge a fraction of the spouse's wage income:

$$b \leq \theta_k(1-\delta)\, k' + \theta_w \cdot w. \tag{12}$$

Borrowing b is a choice variable: entrepreneurs may borrow any amount up to the limit. Since both borrowing and lending occur at rate r, only the net savings position $s \equiv d - b$ matters for the entrepreneur's budget, where $d \geq 0$ denotes deposits and $b \geq 0$ denotes borrowing. The net savings must satisfy $s \geq -\bar{b}(k', w)$, where $\bar{b}$ is the relevant collateral limit above. The wage pledgeability parameter θ_w is central to the amplification mechanism: it allows E2 entrepreneurs to take on higher leverage in normal times, but exposes them to greater deleveraging pressure when spousal unemployment risk rises.

6.5 Idiosyncratic Unemployment Risk

Building on the literature showing that countercyclical unemployment risk is a key driver of aggregate fluctuations through precautionary savings (Krusell and Smith, 1998; Ravn and Sterk, 2017; McKay and Reis, 2016; Broer et al., 2025), all agents in our model face countercyclical unemployment risk. Each period, each worker independently draws an employment status $\varepsilon_t \in \{1, 0\}$ from a Bernoulli distribution with unemployment probability:

$$u(A) = \bar{u} + \chi \cdot \max(0,\ 1 - A). \tag{13}$$

When employed, the worker earns wage w. When unemployed, the worker receives unemployment benefits at replacement rate ξ:

$$\tilde{w}(\varepsilon) = \begin{cases} w - \tau_{\text{ui}} & \text{if } \varepsilon = 1 \text{ (employed)}, \\ \xi \cdot w & \text{if } \varepsilon = 0 \text{ (unemployed)}. \end{cases} \tag{14}$$

The i.i.d. nature of the shock means each period brings a fresh draw, generating precautionary savings demand that rises with $u(A)$. Standalone workers in both economies face the same unemployment risk. In E1, this affects the bond supply through workers' precautionary savings behavior but does not directly enter the entrepreneur's problem (since E1 entrepreneurs have no spouse). In E2, the entrepreneur internalizes the spouse's unemployment risk through the joint household budget constraint.

6.6 Budget Constraints and Entrepreneur Turnover

E1 (Standalone Entrepreneur). The entrepreneur chooses capital k' and net savings s:

$$c + k' + s = a, \qquad c > 0, \quad k' \geq 0, \quad s \geq -\theta_k(1-\delta)\, k'. \tag{15}$$

Next-period wealth is $a' = R^{\text{gross}}(k', z') + (1+r)\, s$, where R^{gross} denotes gross production returns including undepreciated capital.

E2 (Entrepreneur-Worker Household). The household budget is identical in structure but with a larger borrowing limit:[8]

$$c + k' + s = a, \qquad s \geq -\left[\theta_k(1-\delta)\, k' + \theta_w \cdot w\right]. \tag{16}$$

Next-period wealth depends on the employment realization:

$$a'(\varepsilon) = \tilde{w}(\varepsilon) + R^{\text{gross}}(k', z') + (1+r)\, s. \tag{17}$$

[8]Consumption c is measured at the household level. Per-capita comparisons between E1 and E2 require an equivalence scale adjustment, which does not affect optimal policies or the amplification results.

Entrepreneur turnover. Each period, entrepreneurs survive with probability λ and exit with probability $1-\lambda$. Exiting entrepreneurs are replaced by new entrants who inherit a fraction φ of the exiting entrepreneur's wealth, subject to a floor $\underline{a}$: the entry wealth is $\hat{a} = \max(\varphi \cdot a, \underline{a})$. New entrants draw their initial productivity from the ergodic distribution π_z. This turnover structure, following Quadrini (2000), prevents wealth from accumulating indefinitely and generates a realistic wealth distribution.

6.7 Entrepreneur's Recursive Problem

Under Epstein-Zin preferences with risk aversion γ_i and IES ψ, the entrepreneur solves:

$$V(a,z) = \max_{c,\,k',\,s} \left\{ \left[(1-\beta)\, c^{\rho} + \beta\, \mathrm{CE}(a,z)^{\rho} \right]^{1/\rho} \right\}. \tag{18}$$

The certainty equivalent integrates over productivity transitions, employment shocks (in E2), and entrepreneur turnover:

$$\mathrm{CE}(a,z) = \left[\lambda \cdot_{z',\varepsilon} V(a',z')^{1-\gamma_i} + (1-\lambda) \cdot V^{\text{entry}}(a)^{1-\gamma_i} \right]^{1/(1-\gamma_i)}, \tag{19}$$

where $V^{\text{entry}}(a) = \left[\sum_z \pi_z\, V(\hat{a},\, z)^{1-\gamma_i} \right]^{1/(1-\gamma_i)}$ is the certainty equivalent of entering at wealth $\hat{a} = \max(\varphi \cdot a,\, \underline{a})$ with a random productivity draw. The expectation over next-period states is:

$$_{z',\varepsilon} V^{1-\gamma_i} = \sum_{z'} P(z'|z) \sum_{\varepsilon \in \{0,1\}} p(\varepsilon) \cdot V(a'(\varepsilon),\, z')^{1-\gamma_i}, \tag{20}$$

where $p(1) = 1 - u(A)$ and $p(0) = u(A)$. In E1, the ε sum collapses to a single term (no spouse). The entry value is state-dependent: wealthier exiting entrepreneurs spawn wealthier entrants, capturing partial wealth transmission across entrepreneurial generations.

6.8 Worker's Problem

Workers solve a consumption-savings problem with Epstein-Zin preferences, risk aversion γ_W, and shared IES ψ:

$$V_w(b, z_w) = \max_{b' \geq 0} \left\{ \left[(1-\beta)\, c_w^{\rho} + \beta\, \mathrm{CE}_w^{\rho} \right]^{1/\rho} \right\}, \tag{21}$$

subject to $c_w = \tilde{w}(\varepsilon) \cdot z_w + (1+r)\, b - b'$ with $b' \geq 0$, where z_w is idiosyncratic worker productivity following its own AR(1) process. The certainty equivalent integrates over both productivity transitions and next-period employment draws:

$$\mathrm{CE}_w = \left[\sum_{z_w'} P(z_w'|z_w) \sum_{\varepsilon'} p(\varepsilon'; A) \cdot V_w(b', z_w', \varepsilon')^{1-\gamma_W} \right]^{1/(1-\gamma_W)}. \tag{22}$$

Since unemployment is i.i.d., the transition probabilities are independent of the current employment state. Rising unemployment under a negative TFP shock increases precautionary savings demand by all workers, affecting bond supply in both E1 and E2 symmetrically. Standalone workers face the same UI tax and benefit treatment as E2 spouses.

Unemployment insurance. All employed workers pay a balanced-budget UI payroll tax $\tau_{\text{ui}} = u\xi w/(1-u)$. The balance condition $(1-u)\, N_W\, \tau_{\text{ui}} = u\, N_W\, \xi\, w$ ensures that aggregate contributions equal aggregate benefit payments. Entrepreneurs do not pay τ_{ui} (they have no wage income in E1, and only the spouse's wage is subject to the tax in E2).

6.9 Equilibrium

The equilibrium is characterized by an interest rate r^* that clears the unified asset market:

$$S_w + D_E = K_C(r) + B_E, \tag{23}$$

where S_w is aggregate worker savings, D_E and B_E are entrepreneur deposits and borrowing, and $K_C(r)$ is endogenous corporate capital pinned by no-arbitrage. This unified structure follows Cagetti and De Nardi (2006) and **?**: aggregate savings finance corporate capital and

entrepreneurial borrowing simultaneously, and the interest rate r adjusts to clear the market. Since the no-arbitrage condition immediately determines the wage from r, no separate wage-clearing loop is needed.

Labor market. Only employed workers supply labor. Each entrepreneur hires optimal labor $n^*(k, z; w)$ at the equilibrium wage, generating aggregate entrepreneurial labor demand. The competitive corporate sector absorbs the residual:

$$N_C = (1 - u)\, N_W - N_E^{\text{labor}}. \tag{24}$$

Corporate capital then follows as $K_C = (K_C/N_C) \times N_C$. The wage w is pinned by the no-arbitrage condition, which determines the capital-labor ratio and hence w directly from r.

Amplification. The amplification measure compares output responses to the TFP shock:

$$\text{Amplification} = \frac{\Delta Y_{E2}/Y_{E2}^{\text{base}}}{\Delta Y_{E1}/Y_{E1}^{\text{base}}}. \tag{25}$$

A ratio greater than one means the household economy amplifies the shock.

6.10 Solution Method

The equilibrium is computed using Brent's method on the asset market excess demand function $\text{ED}(r) = (K_C + B_E) - (S_w + D_E)$. For each candidate rate r, the wage w is computed directly from the no-arbitrage condition, eliminating the need for an inner wage loop. The entrepreneur VFI solves a nested optimization over capital k (discrete grid) and net savings s (continuous via Brent's bounded minimization, which combines golden-section search with parabolic interpolation). This continuous optimization replaces the discrete grid search over savings and yields a smoother excess demand function. The worker VFI is a standard discrete-grid problem. Stationary distributions are computed by forward iteration on the transition operator until convergence. The four economies are solved sequentially with warm-started interest rate brackets: E1 base uses the full search range, and each subsequent economy uses

a narrow window centered on a previously solved r^*. Full algorithmic details, including pseudocode for each subroutine, are provided in Appendix B (Algorithms 2–7).

Both the entrepreneur and worker productivity grids are constructed via the Tauchen method and normalized so that $\sum_i \pi_i z_i = 1$ (correcting for Jensen's inequality), ensuring that aggregate labor income equals the corporate wage bill. The precise estimation uses $N_a = 80$ entrepreneur wealth grid points (log-spaced), $N_k = 30$ capital grid points, $N_z = 5$ productivity states, and $N_{a,w} = 150$ worker wealth grid points. VFI convergence tolerance is 10^{-3} and asset market clearing tolerance is $|\text{ED}| < 0.01$. The interest rate is searched over $r \in [0.001, 0.040]$. Additional computational parameters are reported in Table B.6 in the appendix.

6.11 Calibration

Table 1 reports the parameter values. Most parameters follow standard values in the literature. The discount factor, capital share, and depreciation rate are set to conventional annual values. The decreasing returns parameter ν follows Cagetti and De Nardi (2006). Entrepreneur risk aversion is heterogeneous, with L-types less risk-averse than H-types, while workers are the most risk-averse. The IES is set below unity, consistent with estimates in the macro-finance literature. The capital collateral fraction θ_k governs borrowing capacity against physical capital; wage pledgeability θ_w determines the additional collateral available in E2. Entrepreneur turnover follows Quadrini (2000). The unemployment parameters are chosen so that baseline unemployment is 5% and a 5% TFP shock doubles it to 10%, consistent with evidence in Ravn and Sterk (2017). The replacement rate implies a moderate benefit relative to the wage.

6.12 Results

Table 2 reports the equilibrium outcomes for both economies under the baseline TFP ($A = 1$) and after a 5% TFP shock ($A = 0.95$). The shock doubles the unemployment rate from 5% to 10%. E2 entrepreneurs operate with substantially higher leverage than E1 counterparts (0.526 vs. 0.299), reflecting the additional collateral from wage pledgeability. The household economy (E2) experiences a larger total output decline (-11.24%) than the standalone economy

Table 1: Parameter Calibration

Parameter	Symbol	Value	Source / Target
Discount factor	β	0.96	Standard
L-type risk aversion	γ_L	2.0	Low risk aversion
H-type risk aversion	γ_H	3.0	High risk aversion
Worker risk aversion	γ_W	4.0	Higher than entrepreneurs
IES	ψ	0.80	Macro-finance literature
Capital share	α	0.33	Standard
Depreciation rate	δ	0.10	Annual
Entrep. scale	θ	0.90	Calibrated
Entrep. DRS	ν	0.88	Cagetti and De Nardi (2006)
Capital collateral	θ_k	0.60	Fraction of capital pledgeable
Wage collateral	θ_w	0.20	E2 only
Survival rate	λ	0.90	Quadrini (2000)
Entry wealth frac.	φ	0.50	Bequest at turnover
Entry wealth floor	$\underline{a}$	0.50	Minimum entry wealth
Prod. persistence	φ_z	0.80	AR(1) persistence
Prod. volatility	σ_z	0.30	Shock std. dev.
Fraction L-type	f_L	0.50	Equal split
Mass entrepreneurs	N_E	0.10	
Mass workers	N_W	0.90	
Baseline unemp. rate	$\bar{u}$	0.05	5% steady state
Unemp. sensitivity	χ	1.00	$u(0.95) = 10\%$
Replacement rate	ξ	0.50	Moderate benefits

(-11.02%), and the entrepreneurial sector contracts much more sharply in E2 (-11.56% vs. -6.10%), yielding an amplification ratio of 1.02.

6.12.1 E2 Replicates Empirical Patterns of Capital and Leverage Response

A key empirical finding in the entrepreneurship literature is that entrepreneurs respond heterogeneously to aggregate shocks: risk-averse entrepreneurs cut capital more aggressively, and leverage falls during downturns. Figure 1 shows the percentage change in average capital and leverage by entrepreneur type after the TFP shock. In E2, both L-types and H-types reduce capital, with H-types cutting substantially more (-15.3% vs. -9.2%). This heterogeneity arises because the unemployment shock compounds the productivity shock asymmetrically: risk-averse H-types, who are already reluctant to bear idiosyncratic risk, respond more aggressively to the increased wealth volatility from spousal unemployment. Leverage falls across the board in E2 (-3.7% on average), consistent with the precautionary deleveraging channel—

Table 2: Baseline Results

	E1 (Standalone)	E2 (Household)
Y baseline	1.5189	1.5052
Y shock ($A = 0.95$)	1.3515	1.3360
$\Delta Y/Y$ (%)	-11.02	-11.24
Y_E/Y baseline (%)	13.4	15.6
K_E baseline	0.1664	0.1845
$\Delta K_E/K_E$ (%)	-4.23	-12.28
r^* baseline	0.0062	0.0088
r^* shock	0.0035	0.0063
Leverage B/K (base)	0.299	0.526
Leverage B/K (shock)	0.299	0.506
u baseline	5.0%	5.0%
u shock	10.0%	10.0%
Amplification	**1.020**	

entrepreneurs internalize the rising probability that their spouse loses income and voluntarily reduce borrowing.

E1 fails to replicate either pattern. Capital reductions are much smaller and nearly uniform across types (-5.5% for L-types, -2.9% for H-types), and leverage is essentially unchanged ($+0.3\%$). Without spousal income in the budget constraint, E1 entrepreneurs face only the direct productivity shock and have no unemployment risk to internalize. The collateral constraint binds similarly across types, producing homogeneous responses that miss the empirical heterogeneity.

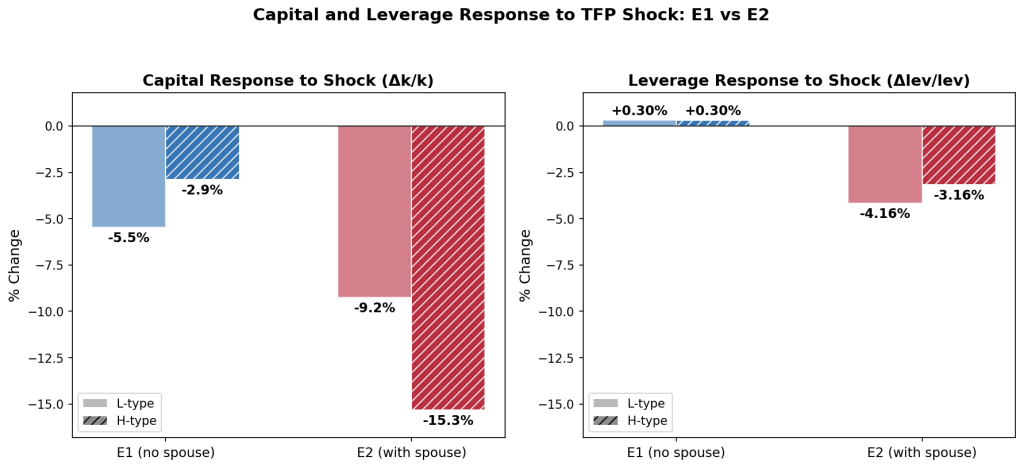


Figure 1: Percentage change in average capital (left) and leverage (right) after a 5% TFP shock, by entrepreneur type. E2 exhibits larger and more heterogeneous capital reductions, and leverage falls—matching empirical patterns. E1 shows small, uniform responses and no deleveraging.

6.12.2 Wealth Concentration and Capital Allocation

Following Moll (2014), we analyze the wealth shares $\omega(z)$—the fraction of total entrepreneurial wealth held by each productivity type—and compare them to the population mass $\psi(z)$. When $\omega(z)/\psi(z)$ is increasing in z, capital is efficiently concentrated on productive entrepreneurs through the self-financing mechanism: persistent productivity allows high-z types to accumulate wealth over time, relaxing their borrowing constraints.

Figure 2 shows that E2 achieves better capital allocation than E1. The capital share of the highest-productivity entrepreneurs is larger in E2, and the productivity-weighted wealth measure $\sum z \cdot \omega(z)$ is higher (1.61 vs. 1.51). The mechanism is twofold. First, wage income in E2 provides a floor that prevents low-productivity entrepreneurs from draining to zero, but this floor also means that the *relative* wealth share of high-z types is amplified by their superior production income. Second, the additional collateral from wage pledgeability allows productive entrepreneurs to scale up faster through borrowing, concentrating capital where it is most productive.

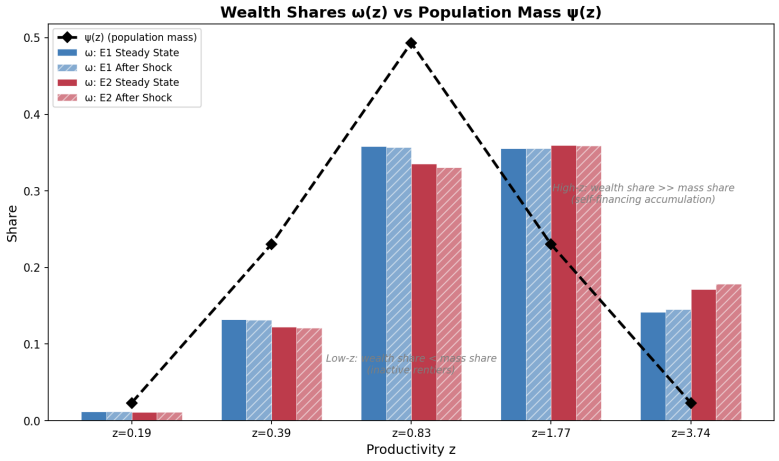


Figure 2: Capital allocation $\kappa(z)$ by productivity type. E2 allocates more capital to high-productivity entrepreneurs than E1, reflecting better self-financing through the wage income channel.

After the TFP shock, wealth shares shift further toward high-z types, and the reallocation is stronger in E2 (Figure 3). At the highest productivity level, the wealth share rises by 3.7% in E2 versus 2.4% in E1. By the Moll TFP measure, this reallocation *improves* allocative efficiency—but the improvement is driven by distress rather than productive accumulation. Middle- and

low-productivity entrepreneurs lose wealth rapidly due to the compound effect of lower profits and spousal unemployment, while wealthy high-z entrepreneurs are partially insulated by larger buffers and the cheap credit from falling interest rates. This creates a fundamental tension: the shock simultaneously reduces total output and improves the composition of remaining wealth.

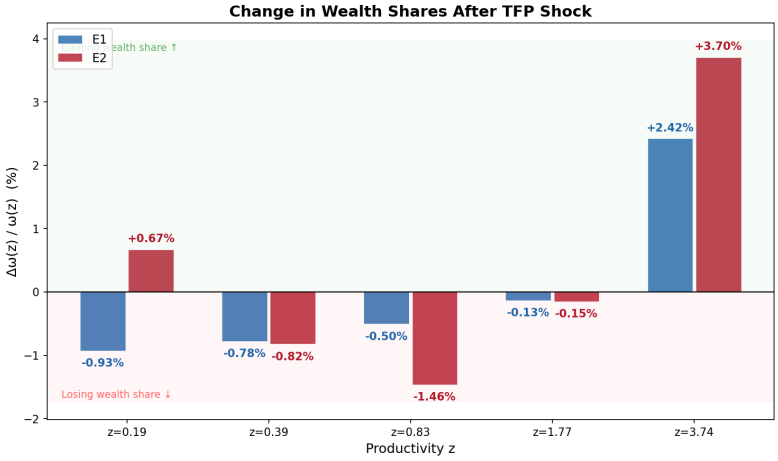


Figure 3: Change in wealth shares $\Delta\omega(z)$ after the TFP shock. E2 exhibits stronger reallocation toward high-productivity types, driven by heterogeneous exposure to spousal unemployment.

6.12.3 Output Decline and Amplification

The total output decline in E2 (-11.24%) exceeds E1 (-11.02%) despite E2 having better capital allocation. The explanation lies in the entrepreneurial sector. E2's entrepreneurial output falls by -11.56% versus only -6.10% in E1. Three reinforcing channels drive this gap. First, the *direct unemployment fear channel*: when unemployment doubles, E2 entrepreneurs face increased variance of next-period wealth from the compound lottery over productivity and employment shocks, triggering precautionary deleveraging. Entrepreneurial capital falls by -12.3% in E2 versus -4.2% in E1, and borrowing drops by -15.5% versus -4.0%. Second, the *extensive margin*: some E2 entrepreneurs whose spouses become unemployed fall below the productivity cutoff for profitable production, exiting the active entrepreneurial sector entirely. Third, the *general equilibrium feedback*: the sharper contraction in E2 entrepreneurial credit demand interacts with workers' rising precautionary savings to determine the equilibrium interest rate, which falls more in E2 (from 0.88% to 0.63%) than in E1 (from 0.62% to 0.35%).

6.12.4 Propagation of Negative Productivity Shocks

A distinctive feature of E2 is that the household structure transforms a transitory aggregate TFP shock into a persistent reallocation of entrepreneurial resources. In E1, the productivity shock reduces output roughly proportionally across the economy, with the entrepreneurial and corporate sectors contracting at similar rates and leverage unchanged. The shock is "clean" in the sense that it does not alter the cross-sectional distribution of entrepreneurial capital or wealth shares in any economically significant way.

In E2, the same TFP shock triggers a cascade of heterogeneous adjustments. The doubling of unemployment creates a binary lottery within each productivity group: entrepreneurs whose spouses retain employment weather the shock with buffers intact, while those whose spouses lose their jobs face simultaneous loss of income and collateral. This heterogeneous exposure reshapes the wealth distribution, pushing a significant mass of entrepreneurs below activation thresholds and concentrating surviving capital on a smaller set of high-productivity, high-wealth operators. Even after the TFP shock reverts, these distributional consequences persist: entrepreneurs who were pushed to low wealth must slowly rebuild through self-financing before they can re-enter production. The precautionary deleveraging is also slow to reverse, as entrepreneurs who experienced spousal unemployment maintain elevated savings buffers. The household economy thus exhibits endogenous persistence—a one-period TFP shock has longer-lasting effects on the distribution of entrepreneurial capital, leverage, and output than the same shock in the standalone economy.

7 Conclusion

Entrepreneurial firms are typically owned and operated by individuals who are part of multi-person households. As a result, business decisions are made in the presence of household income risk that is often uninsurable and closely tied to the owner's personal finances. Despite this, most work on firm behavior abstracts from the household and treats entrepreneurs as standalone decision makers. If household risk cannot be fully separated from business risk, this abstraction may miss an important channel through which economic shocks affect firm policies.

We address this question by exploiting the unexpected collapse in global oil prices in 2014 as a quasi-experimental increase in unemployment risk for households with ties to the oil sector. Using comprehensive Canadian administrative data that link entrepreneurs to their spouses and to detailed firm balance sheets, we compare otherwise similar firms whose owners differ only in their exposure to this household-level shock. We show that entrepreneurs exposed to higher spousal unemployment risk respond by contracting firm investment and leverage while increasing dividend payouts. These adjustments arise even when the spouse remains employed, indicating that entrepreneurs respond to heightened risk rather than realized income losses. The pattern of responses is consistent with precautionary motives: entrepreneurs treat the firm as part of the household portfolio and actively rebalance away from risky business income toward safer and more liquid assets when household background risk increases.

Our findings point to a transmission channel that is largely absent from existing models of private firms and entrepreneurial finance. Entrepreneurs do not make investment and payout decisions in isolation: they internalize household balance-sheet risk and actively adjust firm policies to manage exposure to income uncertainty. As a result, labor-market shocks affecting one household member can propagate into firm behavior through precautionary responses by the entrepreneur. This household-to-firm channel provides a new mechanism through which idiosyncratic employment risk dampen capital accumulation and amplify economic fluctuations.

References

Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, 2012, The network origins of aggregate fluctuations, *Econometrica* 80, 1977–2016.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, 2015, House prices, collateral, and self-employment, *Journal of Financial Economics* 117, 288–306.

Amiti, Mary, and David E Weinstein, 2018, How much do idiosyncratic bank shocks affect investment? evidence from matched bank-firm loan data, *Journal of Political Economy* 126, 525–587.

Arezki, Rabah, and Olivier Blanchard, 2015, The 2014 oil price slump: Seven key questions, *VoxEU, January 13* .

Barrot, Jean-Noël, and Julien Sauvagnat, 2016, Input specificity and the propagation of idiosyncratic shocks in production networks, *The Quarterly Journal of Economics* 131, 1543–1592.

Bayer, Christian, Ralph Luetticke, Lien Pham-Dao, and Volker Tjaden, 2019, Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk, *Econometrica* 87, 255–290.

Bena, Jan, Serdar Dinc, and Isil Erel, 2022, The international propagation of economic downturns through multinational companies: The real economy channel, *Journal of Financial Economics* 146, 277–304.

Bermejo, Vicente J, Miguel A Ferreira, Daniel Wolfenzon, and Rafael Zambrana, 2025, How do cash windfalls affect entrepreneurship? evidence from the spanish christmas lottery, *Journal of Financial and Quantitative Analysis* 60, 2229–2258.

Biermann, Marcus, and Kilian Huber, 2024, Tracing the international transmission of a crisis through multinational firms, *The Journal of Finance* 79, 1789–1829.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten, 2016, Consumption inequality and family labor supply, *American Economic Review* 106, 387–435.

Bonikowska, Aneta, and René Morissette, 2012, Earnings losses of displaced workers with stable labour market attachment: Recent evidence from Canada.

Broer, Tobias, Niels-Jakob Harbo Hansen, Per Krusell, and Erik Öberg, 2025, The unemployment-risk channel in business-cycle fluctuations, *International Economic Review* Forthcoming.

Cagetti, Marco, and Mariacristina De Nardi, 2006, Entrepreneurship, frictions, and wealth, *Journal of Political Economy* 114, 835–870.

Caggese, Andrea, Vicente Cuñat, and Daniel Metzger, 2019, Firing the wrong workers: Financing constraints and labor misallocation, *Journal of Financial Economics* 133, 589–607.

Carroll, Christopher D, Karen E Dynan, and Spencer D Krane, 2003, Unemployment risk and precautionary wealth: Evidence from households' balance sheets, *Review of Economics and Statistics* 85, 586–604.

Carvalho, Vasco M, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi, 2021, Supply chain disruptions: Evidence from the great east japan earthquake, *The Quarterly Journal of Economics* 136, 1255–1321.

Carvalho, Vasco M, and Alireza Tahbaz-Salehi, 2019, Production networks: A primer, *Annual Review of Economics* 11, 635–663.

Catherine, Sylvain, 2022, Countercyclical labor income risk and portfolio choices over the life cycle, *The Review of Financial Studies* 35, 4016–4054.

Charness, Gary, and Uri Gneezy, 2012, Strong evidence for gender differences in risk taking, *Journal of economic behavior & organization* 83, 50–58.

Chodorow-Reich, Gabriel, 2014, The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis, *The Quarterly Journal of Economics* 129, 1–59.

Corradin, Stefano, and Alexander Popov, 2015, House prices, home equity borrowing, and entrepreneurship, *The Review of Financial Studies* 28, 2399–2428.

Cravino, Javier, and Andrei A Levchenko, 2017, Multinational firms and international business cycle transmission, *The Quarterly Journal of Economics* 132, 921–962.

Delavande, Adeline L, Gizem Koşar, and Basit Zafar, 2025, Information spillovers within couples: Evidence from a sequential survey of spouses, Technical report, National Bureau of Economic Research.

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G Wagner, 2011, Individual risk attitudes: Measurement, determinants, and behavioral consequences, *Journal of the european economic association* 9, 522–550.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde, 2018, Global evidence on economic preferences, *The quarterly journal of economics* 133, 1645–1692.

Fremeth, Adam, Guy Holburn, and Margaret Loudermilk, 2016, Energy in canada: A statistical overview, *Policy Brief* .

Giroud, Xavier, and Holger M Mueller, 2019, Firms' internal networks and local economic shocks, *American Economic Review* 109, 3617–3649.

Guner, Nezih, Yuliya A Kulikova, and Arnau Valladares-Esteban, 2025, Does the added worker effect matter?, *Review of Economic Dynamics* 56, 101271.

Hess, Gregory D, 2004, Marriage and consumption insurance: What's love got to do with it?, *Journal of Political Economy* 112, 290–318.

Huang, Hsuan-Hua, Hsing-Wen Han, Kuang-Ta Lo, and Tzu-Ting Yang, 2023, Liquidity constraints, cash windfalls, and entrepreneurship: Evidence from administrative data on lottery winners, *arXiv preprint arXiv:2303.17029* .

Hurst, Erik, and Annamaria Lusardi, 2004, Liquidity constraints, household wealth, and entrepreneurship, *Journal of political Economy* 112, 319–347.

IMF, 2015, *World Economic Outlook, October 2014: Legacies, Clouds, Uncertainties* (International Monetary Fund).

Jenter, Dirk, and Fadi Kanaan, 2015, Ceo turnover and relative performance evaluation, *the Journal of Finance* 70, 2155–2184.

Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina, 2012, Credit supply and monetary policy: Identifying the bank balance-sheet channel with loan applications, *American Economic Review* 102, 2301–2326.

Juelsrud, Ragnar E, and Ella Getz Wold, 2025, The importance of unemployment risk for individual savings, *Management Science* 71, 6746–6769.

Kilian, Lutz, and Christiane Baumeister, 2016, Expecting the unexpected: Why the oil price keeps surprising us, *VoxEU, February 8* .

Krusell, Per, and Anthony A. Smith, 1998, Income and wealth heterogeneity in the macroeconomy, *Journal of Political Economy* 106, 867–896.

Lachowska, Marta, Alexandre Mas, and Stephen A Woodbury, 2020, Sources of displaced workers' long-term earnings losses, *American Economic Review* 110, 3231–3266.

Lucas, Robert E., 1978, On the size distribution of business firms, *Bell Journal of Economics* 9, 508–523.

Mankart, Jochen, and Rigas Oikonomou, 2017, Household search and the aggregate labour market, *The Review of Economic Studies* 84, 1735–1788.

McKay, Alisdair, and Ricardo Reis, 2016, The role of automatic stabilizers in the U.S. business cycle, *Econometrica* 84, 141–194.

Moll, Benjamin, 2014, Productivity losses from financial frictions: Can self-financing undo capital misallocation?, *American Economic Review* 104, 3186–3221.

Morissette, René, Xuelin Zhang, and Marc Frenette, 2007, *Earnings losses of displaced workers: Canadian evidence from a large administrative database on firm closures and mass layoffs* (Statistics Canada Ottawa).

Moskowitz, Tobias J, and Annette Vissing-Jørgensen, 2002, The returns to entrepreneurial

investment: A private equity premium puzzle?, *American Economic Review* 92, 745–778.

OECD, 2015, Back to work: Canada: Improving the re-employment prospects of displaced workers, Technical report, OECD Publishing, Paris.

OECD, 2019, Oecd indicators of employment protection 2019, Technical report, Organisation for Economic Co-operation and Development.

Quadrini, Vincenzo, 2000, Entrepreneurship, saving, and social mobility, *Review of Economic Dynamics* 3, 1–40.

Raposo, Pedro, Pedro Portugal, and Anabela Carneiro, 2021, The sources of the wage losses of displaced workers: the role of the reallocation of workers into firms, matches, and job titles, *Journal of Human Resources* 56, 786–820.

Ravn, Morten O., and Vincent Sterk, 2017, Job uncertainty and deep recessions, *Journal of Monetary Economics* 90, 125–141.

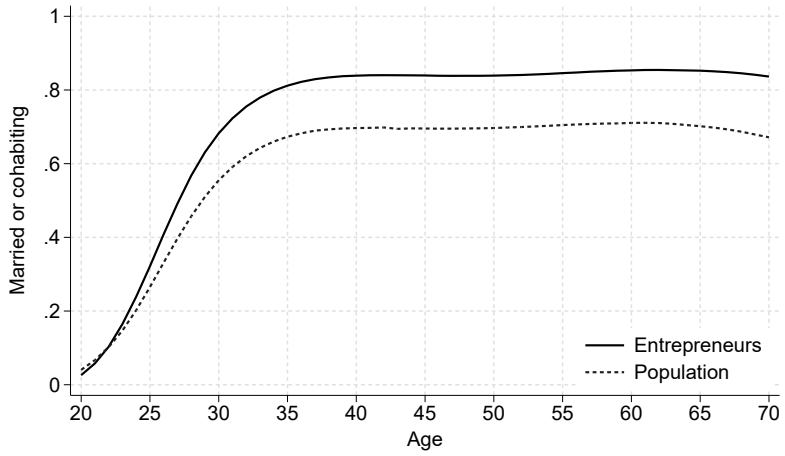
Schmalz, Martin C, David A Sraer, and David Thesmar, 2017, Housing collateral and entrepreneurship, *The Journal of Finance* 72, 99–132.

Stephens, Melvin, Jr, 2002, Worker displacement and the added worker effect, *Journal of Labor Economics* 20, 504–537.

Stepner, Michael, and Allan Garland, 2017, Calipmatch: Stata module for caliper matching without replacement .

Figures

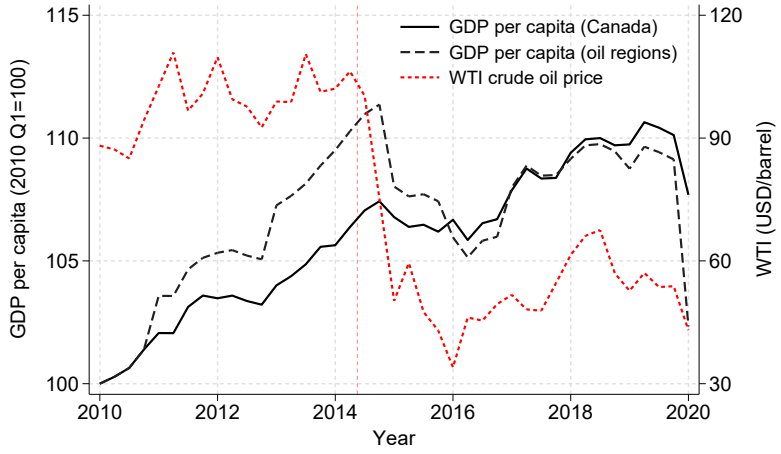
Figure 4: Share of married or cohabiting entrepreneurs



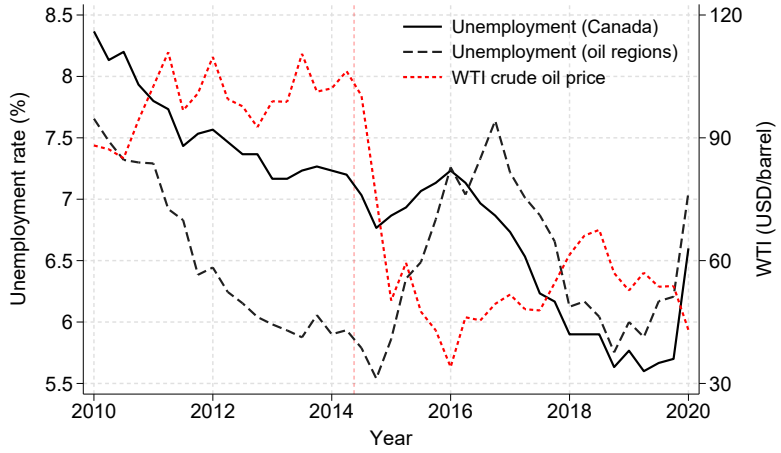
Notes: This figure plots the share of individuals who are married or cohabiting at each age for the universe of entrepreneurs and for the overall population in Canada between 2001 and 2017.

Figure 5: Oil prices and macroeconomic indicators in Canada

(a) Oil price and GDP per capita



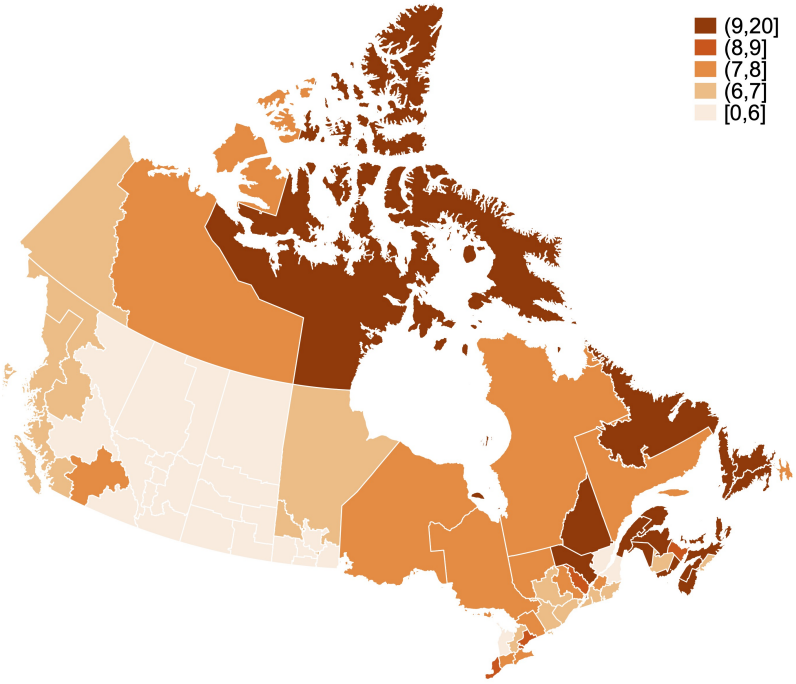
(b) Oil price and unemployment



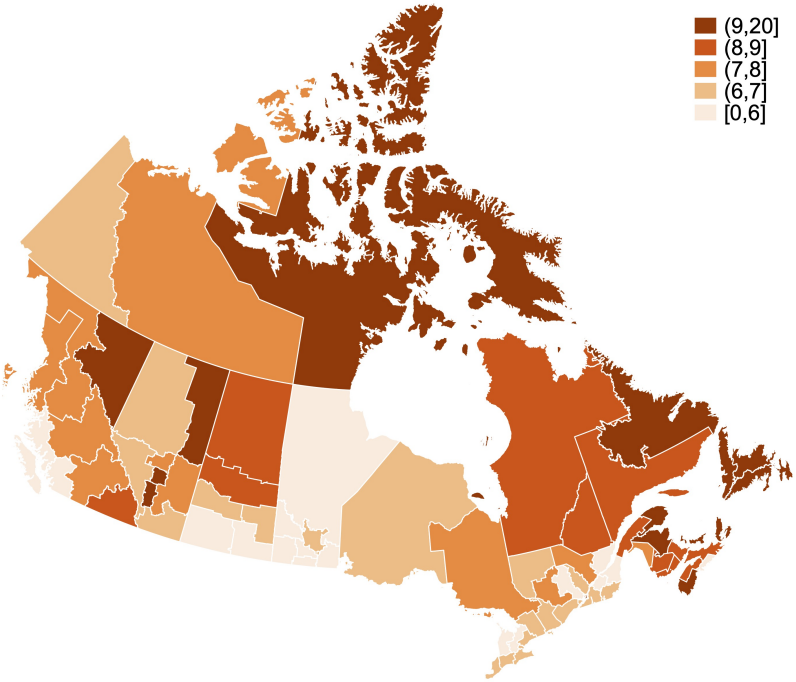
Notes: This figure presents the relationship between oil prices and aggregate macroeconomic conditions in Canada and in the main oil-exposed provinces (Alberta, Saskatchewan, British Columbia, and Newfoundland and Labrador). Panel A plots quarterly WTI crude oil prices (in 2017 USD) against growth in real GDP per capita. National GDP is available at quarterly frequency. Provincial GDP is released annually and converted to quarterly frequency using the Denton interpolation method. Panel B plots oil prices against quarterly unemployment rates. Unemployment and GDP data are sourced from Statistics Canada tables 36-10-0104-01, 36-10-0222-01, 17-10-0009-01, and 14-10-0287-01, .

Figure 6: Unemployment rates across Canadian Economic Regions

(a) Unemployment rates in 2013

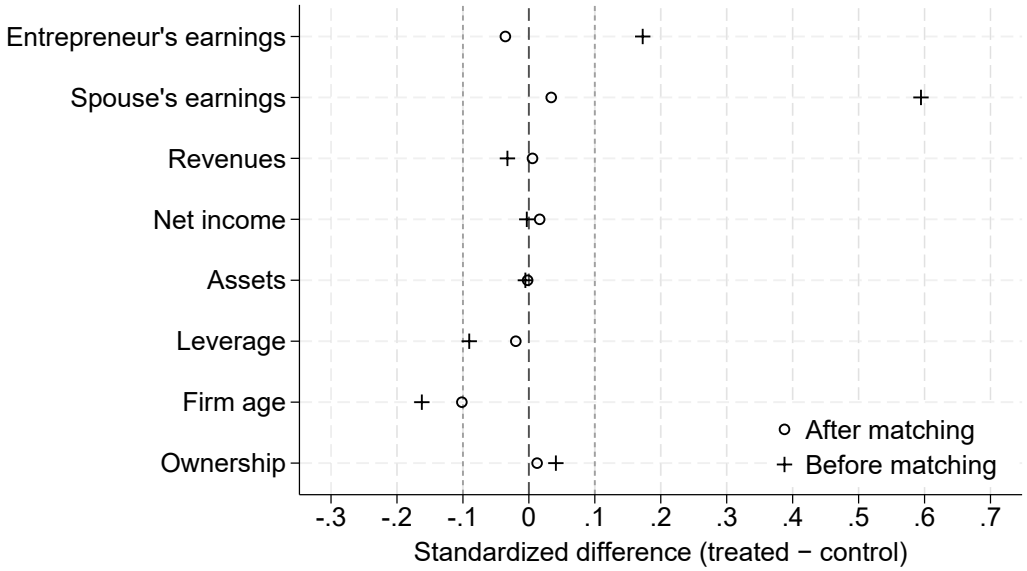


(b) Unemployment rates in 2016



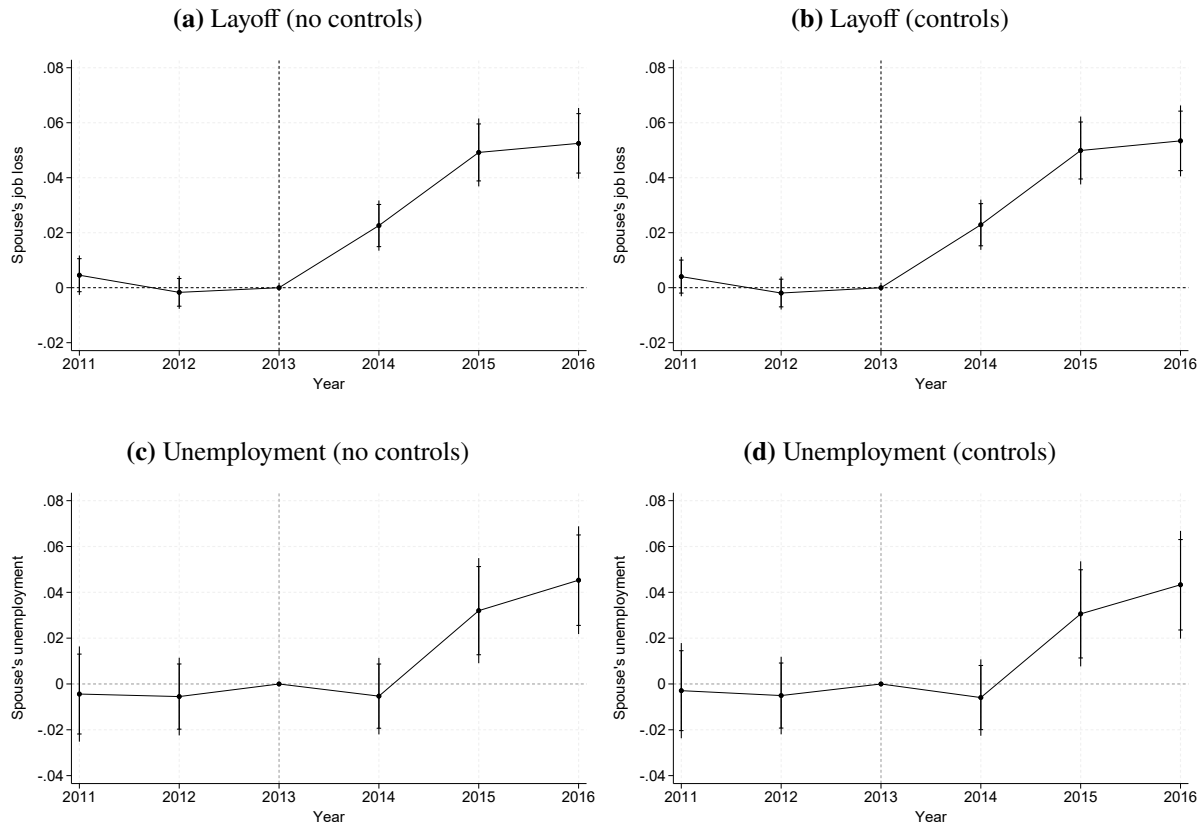
Notes: This figure shows unemployment rates across Canadian Economic Regions in 2013 (Panel A) and 2016 (Panel B). Economic Regions are defined by Statistics Canada as standard geographic units for analysis of regional economic activity. An Economic Region contains multiple commuting zones. Unemployment data are obtained from Statistics Canada table 14-10-0464-01.

Figure 7: Covariate balance



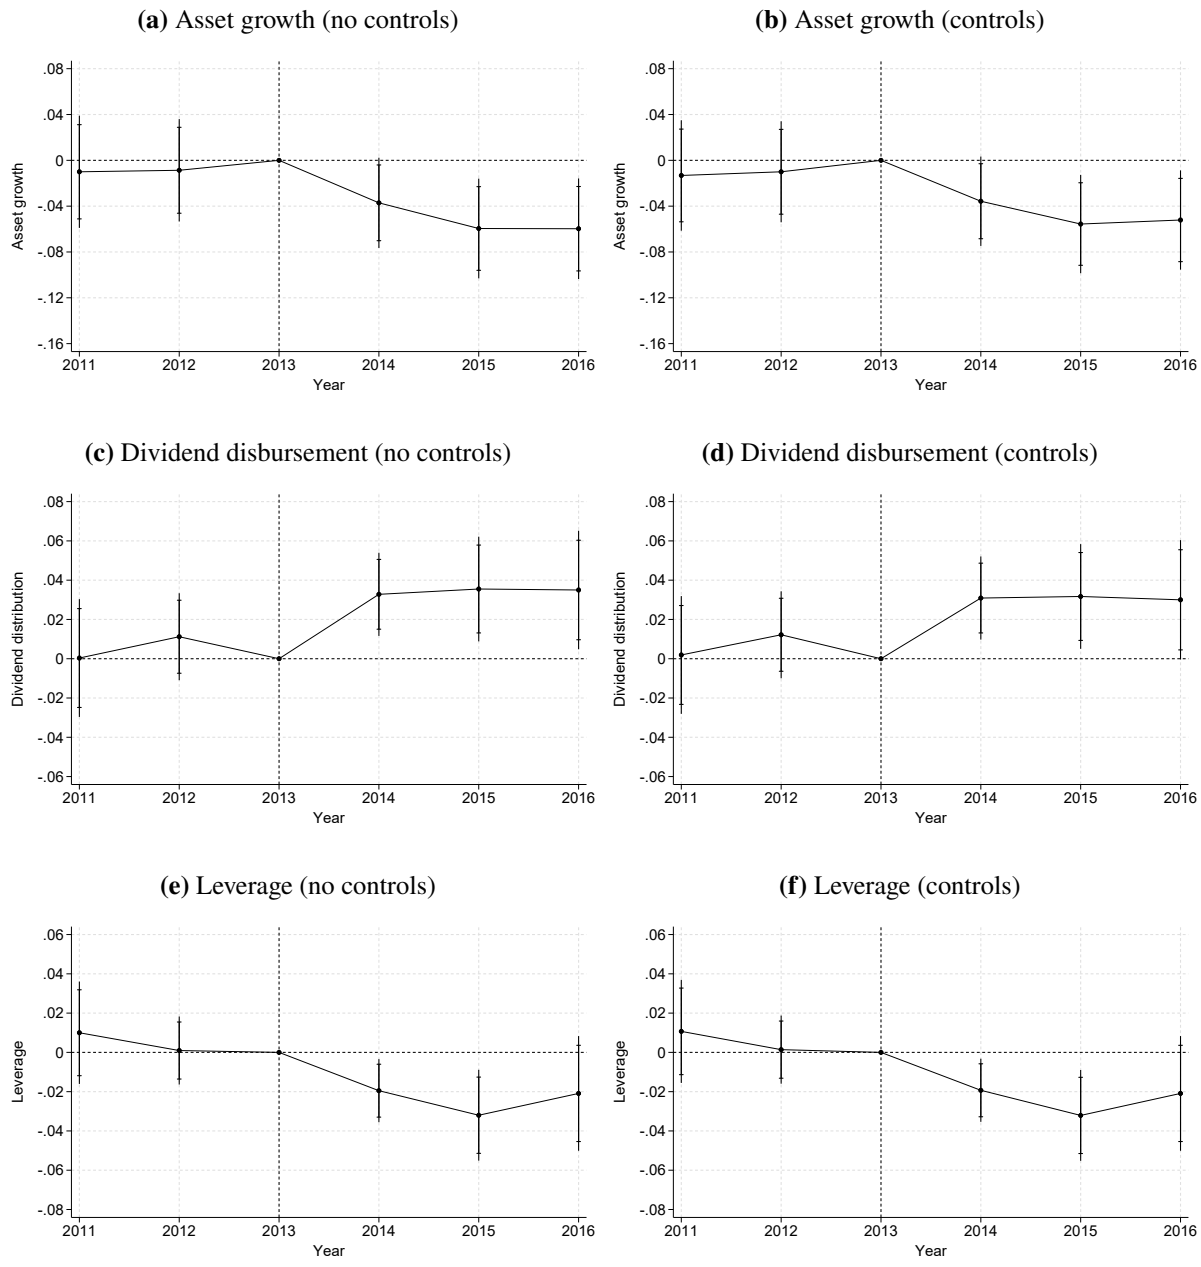
Notes: This figure reports the standardized difference in means between treated and control firms for the continuous variables listed in Table 3. Treated firms are those whose owner's spouse was employed in the oil sector in 2013 and 2014. Potential controls are firms owned by entrepreneurs with spouses employed in other sectors. The standardized difference is computed as $\frac{\bar{X}_T - \bar{X}_C}{\sqrt{\frac{1}{2}\left(s_T^2 + s_C^2\right)}}$, where $\bar{X}_T$ and $\bar{X}_C$ denote sample means and s_T^2 and s_C^2 denote sample variances for treated and control firms, respectively.

Figure 8: Effects of the oil shock on spouse's job loss and unemployment risk



Notes: This figure shows event–study estimates from equation 1 on an indicator indicator for whether an individual experiences a layoff (panels (a)–(b)) or unemployment (panels (c)–(d)). Layoff is an indicator equal to 1 when a worker's employment spell ends and the employer records the reason for the separation as "shortage of work", corresponding to an involuntary, employer-initiated job loss. Unemployment is an indicator equal to 1 if the worker receives unemployment insurance benefits. The treated group consists of spouses employed in the oil sector in both 2013 and 2014, while the matched control group includes spouses employed in other sectors. Control variables include a polynomial for age. Individual effects and commuting zone $\times$ year effects are included. We report 90% and 95% confidence intervals based on standard errors clustered at the individual level.

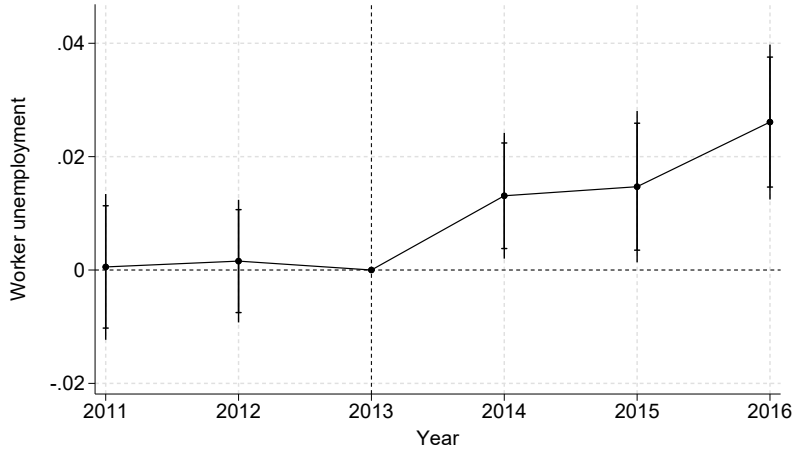
Figure 9: Spouse's unemployment risk and firm outcomes



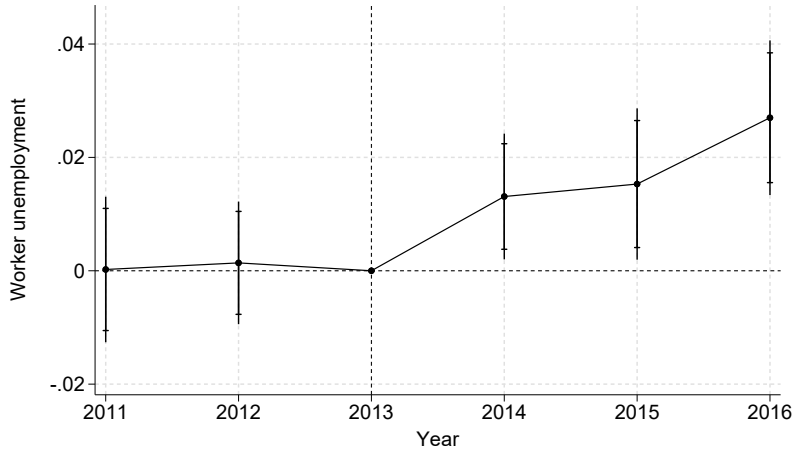
Notes: This figure reports event–study estimates obtained by fitting equation 2 on asset growth, dividend disbursement, and financial leverage. The treated group consists of entrepreneurs with spouses employed in the oil sector in both 2013 and 2014, while the matched control group includes entrepreneurs with spouses employed in other sectors. Asset growth is $\Delta a_t = log(a_t) - log(a_{t-1})$. Dividend disbursement is an indicator indicator for whether an entrepreneur declares dividends in year t. Financial leverage is the ratio of total long term liabilities divided by total assets. Controls include polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of owners. Firm fixed effects and industry $\times$ commuting zone $\times$ year fixed effects are included. We report 90% and 95% confidence intervals based on standard errors clustered at the firm level.

Figure 10: Propagation of household unemployment risk to employees

(a) Worker unemployment (no controls)



(b) Worker unemployment (controls)



Notes: This figure shows event–study estimates from equation 3 on an indicator indicator for whether an worker experiences unemployment. The treated group consists of workers who were employed in both 2013 and 2014 by the entrepreneurs in the treated sample (those whose spouse worked in the oil sector in the same period). The control group consists of workers who are employed by entrepreneurs in the control group. Unemployment is an indicator equal to 1 if the worker receives unemployment insurance benefits. Control variables include a polynomial for age. Individual effects and year effects are included. We report 90% and 95% confidence intervals based on standard errors clustered at the individual level.

Tables

Table 3: Summary statistics

	(1) Matched treated	(2) Matched control	(3) All treated	(4) Potential control
Entrepreneur's earnings (000)	98.84	102.19	97.55	83.71
	(89.84)	(97.94)	(85.26)	(74.74)
Spouse's earnings (000)	106.88	104.07	124.51	62.03
	(82.29)	(83.59)	(131.09)	(69.97)
Couple's wealth (000)	1,720.26	1,789.51	1,295.18	?
	(1,993.72)	(2,390.91)	(1,964.73)	?
Revenue (000)	648.92	638.278	895.51	1,011.88
	(2,067.02)	(1,818.48)	(3,633.31)	(3,530.17)
Net Income (000)	94.24	84.88	132.72	166.76
	(555.62)	(588.63)	(1,412.00)	(15,700.00)
Assets (000)	648.58	652.52	797.27	813.66
	(2,251.87)	(1,853.58)	(3,194.19)	(3,036.40)
Leverage	0.27	0.28	0.28	0.33
	(0.50)	(0.51)	(0.53)	(0.55)
Firm age	10.23	11.17	10.27	11.91
	(8.91)	(9.58)	(9.41)	(10.75)
Ownership (%)	71.54	71.10	68.74	67.28
	(35.20)	(35.12)	(36.16)	(35.11)
Female	0.39	0.39	0.47	0.18
Province:				
Alberta	0.61	0.61	0.60	0.11
British Columbia	0.08	0.08	0.08	0.12
Ontario	0.14	0.14	0.13	0.35
Newfoundland and Labrador	0.01	0.01	0.01	0.01
Quebec	0.12	0.12	0.11	0.32
Saskatchewan	0.02	0.02	0.04	0.02
Other	0.02	0.02	0.01	0.02
Observations	2,365	2,365	3,440	208,320

Notes: Column (1) reports summary statistics for firms whose owner's spouse was employed in the oil sector in 2013–2014 and for whom we can find a matched control entrepreneur-firm pair. Summary statistics are measured in 2013, the year prior to the oil price collapse. Potential controls are firms whose owner's spouse was employed in a non-oil sector during the same period. Matching is performed in 2013 using a caliper matching algorithm (Stepner and Garland, 2017). Each treated firm is matched exactly on 4-digit industry, city, and entrepreneur's sex and within a caliper of 25 log points on the spouse's earnings. Column (2) reports the characteristics of matched treated firms. Column (3) reports the characteristics of all treated firms prior to matching. Column (4) reports the characteristics of all potential control firms. All dollar values are expressed in thousands of 2013 Canadian dollars. Ownership refers to the entrepreneur's equity share in the firm. Standard deviations are reported in parentheses. The number of observations is rounded to the nearest five to comply with Statistics Canada confidentiality requirements.

Table 4: Effects of the oil shock on spouse's labor market outcomes

Panel A. Job loss

	(1)	(2)	(3)	(4)	(5)	(6)
Post × Oil_i	0.043***	0.042***	0.042***	0.055***	0.055***	0.054***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	No	No	Yes	No	No
Province × year FE	No	Yes	No	No	Yes	No
Commuting zone × year FE	No	No	Yes	No	No	Yes
Observations	31,616	31,605	31,415	22,370	22,370	22,240

Panel B. Unemployment

	(1)	(2)	(3)	(4)	(5)	(6)
Post × Oil_i	0.019**	0.019**	0.018**	0.022**	0.022**	0.020**
	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)	(0.009)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	No	No	Yes	No	No
Province × year FE	No	Yes	No	No	Yes	No
Commuting zone × year FE	No	No	Yes	No	No	Yes
Observations	31,615	31,605	31,415	22,370	22,370	22,240

Panel C. Log earnings

	(1)	(2)	(3)	(4)	(5)	(6)
Post × Oil_i	-0.049**	-0.047**	-0.046**	-0.056**	-0.056**	-0.054**
	(0.018)	(0.018)	(0.018)	(0.022)	(0.021)	(0.021)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	No	No	Yes	No	No
Province × year FE	No	Yes	No	No	Yes	No
Commuting zone × year FE	No	No	Yes	No	No	Yes
Observations	31,555	31,545	31,355	22,325	22,325	22,195

Notes: This table reports average treatment effects from equation 1 on an indicator for job loss (Panel A), an indicator for unemployment (Panel B), and log annual earnings (Panel C). The sample covers the years 2010 to 2016. Post is an indicator equal to 1 in 2014 or after. Job loss is equal to 1 when a worker's employment spell ends and the employer records the reason for the separation as "shortage of work", corresponding to an involuntary, employer-initiated job loss. Unemployment is equal to 1 if the worker receives unemployment insurance benefits. Columns (1)–(3) include the full sample. Columns (4)–(6) restrict the sample to oil-exposed provinces (Alberta, Saskatchewan, British Columbia, and Newfoundland and Labrador). The treated group consists of spouses employed in the oil sector in both 2013 and 2014, while the matched control group includes spouses employed in other sectors. Controls include a polynomial in age. Standard errors are reported in parentheses and are clustered at the individual level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. The numbers of observations are rounded to the nearest five to comply with the confidentiality requirements of Statistics Canada.

Table 5: Spouse's unemployment risk and firm outcomes

Panel A. Asset growth

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post × $\text{Oil}_{i(j)}$	-0.042***	-0.034**	-0.043***	-0.035***	-0.050***	-0.040***	-0.059***	-0.054***
	(0.014)	(0.013)	(0.014)	(0.013)	(0.014)	(0.014)	(0.017)	(0.016)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year FE	Yes	Yes	Yes	Yes	No	No	No	No
Commuting zone × year FE	No	No	Yes	Yes	No	No	No	No
Industry × commuting zone × year FE	No	No	No	No	Yes	Yes	Yes	Yes
Observations	25,725	25,725	25,605	25,600	24,020	24,015	17,350	17,345
R^2	0.31	0.33	0.33	0.35	0.41	0.43	0.38	0.40

Panel B. Dividend disbursement

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post × $\text{Oil}_{i(j)}$	0.027**	0.024**	0.028***	0.026**	0.030***	0.029***	0.040***	0.036***
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.013)	(0.013)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year FE	Yes	Yes	Yes	Yes	No	No	No	No
Commuting zone × year FE	No	No	Yes	Yes	No	No	No	No
Industry × commuting zone × year FE	No	No	No	No	Yes	Yes	Yes	Yes
Observations	27,660	27,655	27,340	27,335	25,945	25,935	18,635	18,625
R^2	0.70	0.70	0.72	0.72	0.75	0.75	0.74	0.74

Table 5 (continued)

Panel C. Leverage

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post × $\text{Oil}_{i(j)}$	-0.029***	-0.027**	-0.031***	-0.029**	-0.028***	-0.027**	-0.039***	-0.036***
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.013)	(0.013)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year FE	Yes	Yes	Yes	Yes	No	No	No	No
Commuting zone × year FE	No	No	Yes	Yes	No	No	No	No
Industry × commuting zone × year FE	No	No	No	No	Yes	Yes	Yes	Yes
Observations	26,975	26,975	26,690	26,685	25,250	25,245	18,140	18,130
R^2	0.79	0.79	0.80	0.80	0.83	0.83	0.81	0.81

Notes: This table reports average treatment effects from equation 2 on asset growth (Panel A), dividend disbursement (Panel B), and financial leverage (Panel C). The sample covers the years 2010 to 2016. Post is equal to 1 in 2014 or after. $\text{Oil}_{i(j)}$ is an indicator equal to 1 for entrepreneurs with spouses employed in the oil sector in both 2013 and 2014, while the matched control group includes entrepreneurs with spouses employed in other sectors. Asset growth is defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. Dividend disbursement is an indicator for whether the entrepreneur declares dividends in year t. Financial leverage is the ratio of long-term liabilities to total assets. Columns (1)–(6) use the full sample of incorporated firms, while columns (7)–(8) restrict the sample to oil-exposed provinces (Alberta, Saskatchewan, British Columbia, and Newfoundland and Labrador). The treated group consists of entrepreneurs whose spouses were employed in the oil sector in both 2013 and 2014, and the matched control group includes entrepreneurs whose spouses worked in other sectors. Controls include polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of owners. Standard errors are reported in parentheses and are clustered at the firm level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. The numbers of observations are rounded to the nearest five to comply with the confidentiality requirements of Statistics Canada.

Table 6: Operating outcomes

	Revenue growth		**Expense growth**		**Profit margin**		**Productivity**		**Survival**	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Post $\times$ $\text{Oil}_{i(j)}$	$-^{*}$	$-^{*}$	$-$	$-$	$-^{*}$	$-^{*}$	$-^{*}$	$-^{*}$	$-^{*}$	$-^{*}$
Controls	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry $\times$ CS $\times$ year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations										
R^2										

Notes: **These results await release from Statistics Canada.** This table reports average treatment effects from equation 2 on revenue growth, expense growth, profit margin, productivity, and survival. The sample covers the years 2010 to 2016. Post is equal to 1 in 2014 or after. $\text{Oil}_{i(j)}$ is an indicator equal to 1 for entrepreneurs with spouses employed in the oil sector in both 2013 and 2014. The matched control group includes entrepreneurs with spouses employed in other sectors. Controls include polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of owners. Standard errors are reported in parentheses and are clustered at the firm level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. The numbers of observations are rounded to the nearest five to comply with the confidentiality requirements of Statistics Canada.

Table 7: Propagation of household unemployment risk to employees

	(1)	(2)	(3)	(4)	(5)	(6)
Post $\times$ $\text{Oil}_{i(l)}$	0.017***	0.016***	0.014***	0.014***	0.010**	0.010**
	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
Post $\times$ Young					0.098***	0.010
					(0.007)	(0.0073
Post $\times$ Young $\times$ $\text{Oil}_{i(l)}$					0.019**	0.020**
					(0.009)	(0.001)
Controls	No	Yes	No	Yes	No	Yes
Worker FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	No	No	No	No
Commuting zone $\times$ year FE	No	No	Yes	Yes	Yes	Yes
Observations	204,805	204,805	204,715	204,715	204,715	204,715

Notes: This table reports average trearment effects from equation 3 on an indicator indicator for whether an worker experiences unemployment. Post is equal to 1 in 2014 or after. $\text{Oil}_{i(l)}$ is an indicator equal to 1 if worker l was employed in both 2013 and 2014 by an entrepreneur whose spouse worked in the oil sector in the same period. Workers employed by these entrepreneurs form the treated group. The control group consists of workers employed by entrepreneurs whose spouses worked outside the oil sector. Young is equal to 1 if a worker is younger than 30 years old. Unemployment is an indicator equal to 1 if the worker receives unemployment insurance benefits. Control variables include a polynomial for age. Standard errors are reported in parentheses and are clustered at the individual level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. The numbers of observations are rounded to the nearest five to comply with the confidentiality requirements of Statistics Canada.

Table 8: Heterogeneous effects on investment by entrepreneur's risk aversion

	(1)	(2)	(3)	(4)	(5)	(6)
Post $\times$ $\text{Oil}_{i(j)}$	-0.036^{**}	-0.031^{**}	-0.041^{***}	-0.035^{**}	-0.043^{***}	-0.036^{**}
	(0.016)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Post $\times$ $\text{Oil}_{i(j)}$ $\times$ $\mathbb{I}\{\text{Age} > \text{p75}\}$	-0.060^{*}	-0.058^{*}				
	(0.035)	(0.034)				
Post $\times$ $\text{Oil}_{i(j)}$ $\times$ $\mathbb{I}\{\text{Female, Age} > \text{p75}\}$			-0.080^{*}	-0.083^{**}		
			(0.042)	(0.041)		
Post $\times$ $\text{Oil}_{i(j)}$ $\times$ $\mathbb{I}\{\text{Female, Age} > \text{p75, High bargaining power}\}$					-0.124^{*}	-0.132^{**}
					(0.064)	(0.065)
Controls	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry $\times$ commuting zone $\times$ year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	24,020	24,015	24,020	24,015	24,020	24,015
R^2	0.409	0.425	0.409	0.425	0.409	0.425

Notes: This table reports heterogeneous treatment effects estimated from equation 2 on asset growth. The sample covers the years 2010 to 2016. Post is equal to 1 in 2014 or after. $\text{Oil}_{i(j)}$ is an indicator equal to 1 for entrepreneurs with spouses employed in the oil sector in both 2013 and 2014, while the matched control group includes entrepreneurs with spouses employed in other sectors. $\mathbb{I}\{\text{Age} > \text{p75}\}$ is an indicator equal to 1 if the entrepreneur is older than the 75th percentile. $\mathbb{I}\{\text{Female, Age} > \text{p75}\}$ is an indicator equal to 1 if the entrepreneur is older than the 75th percentile and a woman. $\mathbb{I}\{\text{Female, Age} > \text{p75, High bargaining power}\}$ is an indicator equal to 1 if the entrepreneur is older than 55, a woman, and her share of household wealth is above the sample median of the distribution of relative wealth within couples. Wealth is proxied by the sum of income from all sources since 2001. We interpret higher relative wealth as a proxy for greater intra-household bargaining power. The dependent variable is asset growth, defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. The treated group consists of entrepreneurs whose spouses were employed in the oil sector in both 2013 and 2014, and the matched control group includes entrepreneurs whose spouses worked in other sectors. Controls include polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of owners. Standard errors are reported in parentheses and are clustered at the firm level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. The numbers of observations are rounded to the nearest five to comply with the confidentiality requirements of Statistics Canada.

Table 9: Heterogeneous effects on investment by industry risk

	Low risk				High risk			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post × $\text{Oil}_{i(j)}$	−0.013	−0.012	−0.024	−0.021	−0.066***	−0.056**	−0.064***	−0.0532**
	(0.017)	(0.017)	(0.018)	(0.018)	(0.022)	(0.022)	(0.023)	(0.023)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year FE	Yes	Yes	Yes	Yes	No	No	No	No
Commuting zone × year FE	No	No	Yes	Yes	No	No	No	No
Industry × commuting zone × year FE	No	No	No	No	Yes	Yes	Yes	Yes
Observations	13,200	13,195	12,010	12,005	11,850	11,850	11,140	11,140
R^2	0.372	0.384	0.452	0.462	0.347	0.369	0.387	0.408

Notes: This table reports heterogeneous treatment effects estimated from equation 2 on asset growth. The sample covers the years 2010 to 2016. Post is equal to 1 in 2014 or after. $\text{Oil}_{i(j)}$ is an indicator equal to 1 for entrepreneurs with spouses employed in the oil sector in both 2013 and 2014, while the matched control group includes entrepreneurs with spouses employed in other sectors. Industries are classified as low- or high-risk based on industry-average ROA volatility, where $\sigma(ROA)$ is computed within firms and then averaged across firms in each industry. The classification uses a median split across industries. The dependent variable is asset growth, defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. The treated group consists of entrepreneurs whose spouses were employed in the oil sector in both 2013 and 2014, and the matched control group includes entrepreneurs whose spouses worked in other sectors. Controls include polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of owners. Standard errors are reported in parentheses and are clustered at the firm level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. The numbers of observations are rounded to the nearest five to comply with the confidentiality requirements of Statistics Canada.

Table 10: Heterogeneous effects on investment by spouse's entrenchment

Panel A. Spouse's job loss

	(1)	(2)	(3)	(4)
Post $\times$ Oil$_i$	0.050^{***}	0.051^{***}	0.050^{***}	0.051^{***}
	(0.006)	(0.006)	(0.006)	(0.006)
Post $\times$ Oil$_i$ $\times$ Manager$_i$	-0.010	-0.010		
	(0.016)	(0.016)		
Post $\times$ Oil$_i$ $\times$ Entrenched manager$_i$			-0.014	-0.014
			(0.015)	(0.016)
Controls	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	No	Yes	No
Commuting zone $\times$ year FE	No	Yes	No	Yes
Observations	29,790	29,585	29,790	29,585
R^2	0.417	0.446	0.417	0.446

Panel B. Asset growth

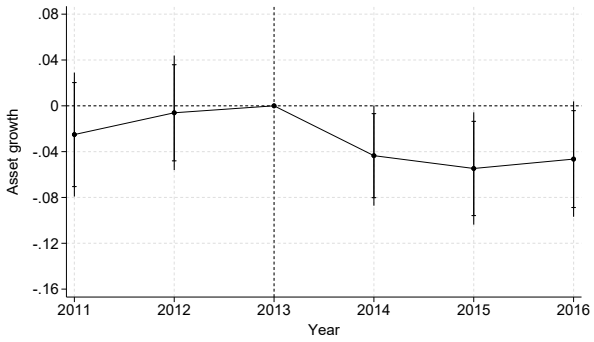
	(1)	(2)	(3)	(4)
Post $\times$ Oil$_{i(j)}$	-0.050^{***}	-0.040^{**}	-0.052^{***}	-0.041^{**}
	(0.016)	(0.016)	(0.016)	(0.016)
Post $\times$ Oil$_{i(j)}$ $\times$ Manager$_{i(j)}$	0.005	0.004		
	(0.045)	(0.044)		
Post $\times$ Oil$_{i(j)}$ $\times$ Entrenched manager$_{i(j)}$			0.011	0.014
			(0.046)	(0.046)
Controls	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry $\times$ commuting zone $\times$ year FE	Yes	Yes	Yes	Yes
Observations	24,020	24,015	24,020	24,015
R^2	0.409	0.425	0.409	0.425

Notes: Panel A reports average treatment effects from equation 1 on an indicator for job loss. The sample covers the years 2010 to 2016. Post is an indicator equal to 1 in 2014 or after. Oil$_i$ is an indicator equal to 1 for spouses employed in the oil sector in both 2013 and 2014. The matched control group includes entrepreneurs with spouses employed in other sectors. Manager$_i$ is an indicator equal to 1 for spouses who are in the top 5th percentile of their firms' earnings distribution. Entrenched manager$_i$ is an indicator equal to 1 for spouses who are in the top 5th percentile of their firms' earnings distribution and above median tenure. Controls include a polynomial in age. Standard errors are reported in parentheses and are clustered at the individual level. Panel B reports heterogeneous treatment effects estimated from equation 2 on asset growth. The sample covers the years 2010 to 2016. The dependent variable is asset growth, defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. The treated group consists of entrepreneurs whose spouses were employed in the oil sector in both 2013 and 2014, and the matched control group includes entrepreneurs whose spouses worked in other sectors. Controls include polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of owners. Standard errors are reported in parentheses and are clustered at the firm level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. The numbers of observations are rounded to the nearest five to comply with the confidentiality requirements of Statistics Canada.

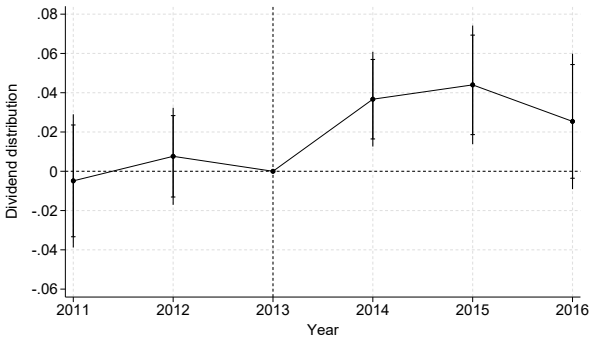
A Additional Figures and Tables

Figure A.1: Firm policies with no realized unemployment

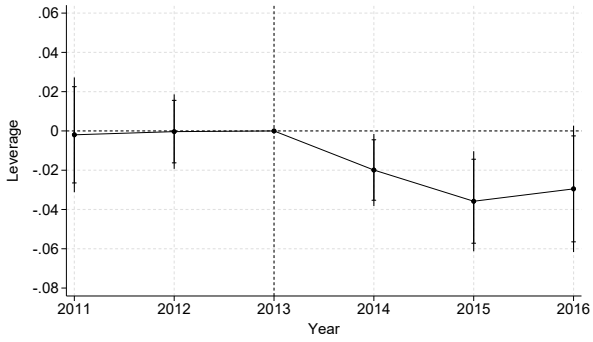
(a) Asset growh



(b) Dividend disbursement



(c) Leverage



Notes: This figure reports event–study estimates obtained by fitting equation 2 on an indicator indicator for whether an entrepreneur declares dividends in year t, conditional on not the spouse not experiencing job loss. The treated group consists of entrepreneurs with spouses employed in the oil sector in both 2013 and 2014, while the matched control group includes entrepreneurs with spouses employed in other sectors. Asset growth is $\Delta a_t = log(a_t) - log(a_{t-1})$. Dividend disbursement is an indicator indicator for whether an entrepreneur declares dividends in year t. Financial leverage is the ratio of total long term liabilities divided by total assets. Controls include polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of owners. Firm fixed effects and industry $\times$ municipality $\times$ year fixed effects are included. We report 90% and 95% confidence intervals based on standard errors clustered at the individual level.

Table A.1: Heterogeneous effects on investment by firm age

	Young firms				Older firms			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post $\times$ Oil$_{i(j)}$	-0.083^{**}	-0.071^{*}	-0.102^{**}	-0.089^{*}	-0.029^{**}	-0.028^{**}	-0.036^{**}	-0.032^{**}
	(0.042)	(0.040)	(0.050)	(0.047)	(0.013)	(0.013)	(0.015)	(0.014)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry $\times$ year FE	Yes	Yes	Yes	Yes	No	No	No	No
Commuting zone $\times$ year FE	No	No	Yes	Yes	No	No	No	No
Industry $\times$ commuting zone $\times$ year FE	No	No	No	No	Yes	Yes	Yes	Yes
Observations	4,755	4,750	3,745	3,740	20,455	20,455	18,575	18,575
R^2	0.510	0.554	0.567	0.608	0.256	0.297	0.365	0.370

Notes: This table reports heterogeneous treatment effects estimated from equation 2 on asset growth. The sample covers the years 2010 to 2016. Post is equal to 1 in 2014 or after. Oil$_{i(j)}$ is an indicator equal to 1 for entrepreneurs with spouses employed in the oil sector in both 2013 and 2014, while the matched control group includes entrepreneurs with spouses employed in other sectors. Young firms are 5 years old or younger in 2014. Older firms are 6 years old or older in 2014. The dependent variable is asset growth, defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. The treated group consists of entrepreneurs whose spouses were employed in the oil sector in both 2013 and 2014, and the matched control group includes entrepreneurs whose spouses worked in other sectors. Controls include polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of owners. Standard errors are reported in parentheses and are clustered at the firm level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. The numbers of observations are rounded to the nearest five to comply with the confidentiality requirements of Statistics Canada.

Table A.2: Robustness to finer geographic effects: postal codes

	(1)	(2)	(3)	(4)
Post $\times$ $\text{Oil}_{i(j)}$	-0.126^{**}	-0.126^{**}	-0.142^{**}	-0.131^{*}
	(0.062)	(0.062)	(0.072)	(0.069)
Controls	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry $\times$ year FE	No	No	Yes	Yes
Postal code $\times$ year FE	Yes	Yes	Yes	Yes
Observations	3,635	3,635	3,000	3,000
R^2	0.658	0.668	0.706	0.716

Notes: This table reports average treatment effects from equation 2 on asset growth using postal codes as the unit of geographic variation. Canadian postal codes are six-character identifiers that represent very fine geographic areas. In urban settings, a postal code typically corresponds to a block-face (one side of a city street between consecutive intersections) or even a single large building. In the administrative data, these postal codes are fully anonymized. Their true geographic locations and boundaries are not provided. Therefore, contiguity across codes is unknown and they cannot be reliably aggregated to larger geographic units. Industry is a 3-digit NAICS code. The sample covers the years 2010 to 2016. Post is equal to 1 in 2014 or after. $\text{Oil}_{i(j}$ is an indicator equal to 1 for entrepreneurs with spouses employed in the oil sector in both 2013 and 2014. The matched control group includes entrepreneurs with spouses employed in other sectors. Postal codes are full The dependent variable is asset growth, defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. Controls include polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of owners. Standard errors are reported in parentheses and are clustered at the firm level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. The numbers of observations are rounded to the nearest five to comply with the confidentiality requirements of Statistics Canada.

Table A.3: Placebo test using spouses who leave the oil industry before the shock

	Spouse's job loss			Asset growth		
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Oil leaver$_i$	-0.002	-0.004	-0.002			
	(0.039)	(0.040)	(0.040)			
Post × Oil leaver$_{i(j)}$				-0.009	0.0140	0.028
				(0.047)	(0.049)	(0.053)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	No	No	No
Firm FE	No	No	No	Yes	Yes	Yes
Year FE	Yes	No	No	No	No	No
Province × year	No	Yes	No	No	No	No
Industry × year FE	No	No	No	Yes	Yes	No
Commuting zone × year FE	No	No	Yes	No	Yes	No
Industry × commuting zone × year FE	No	No	No	No	No	Yes
Observations	3,130	3,120	2,995	2,680	2,625	2,545
R^2	0.403	0.412	0.439	0.372	0.419	0.457

Notes: This table reports results from a placebo test using spouses who leave the oil industry just before the oil shock. Columns (1)-(3) report average treatment effects for spouses based on equation 1. The dependent variable is an indicator for unemployment. Columns (4)-(6) report average treatment effects for entrepreneurial firms based on equation 2. The outcome is asset growth, defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. The sample covers the years 2010 to 2016. Post is equal to 1 in 2014 or after. Oil leaver$_{i(j}$ is an indicator equal to 1 for entrepreneurs with spouses who leave the oil sector in the first half of 2014. The treated group consists of entrepreneurs whose spouses were employed in the oil sector in both 2013 and 2014 but leave the industry between January and June 2014 (just before the collapse in oil prices) for a different sector. The matched control group includes entrepreneurs whose spouses worked in other sectors in both 2013 and 2014 and also leave their sector in the first half of 2014. As in the main sample, treated and control observations are matched based on 2013 characteristics (commuting zone, industry, spouse earnings, and sex). Controls include polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of owners. Standard errors are reported in parentheses and are clustered at the individual level in columns (1)-(3) and at the firm level in columns (4)-(6). *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. The numbers of observations are rounded to the nearest five to comply with the confidentiality requirements of Statistics Canada.

Table A.4: Heterogeneous effects on investment by information advantage

	(1)	(2)	(3)	(4)	(5)	(6)
Post × $\text{Oil}_{i(j)}$	-0.041***	-0.033**	-0.042***	-0.034**	-0.049***	-0.040***
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Post × $\text{Oil}_{i(j)}$ × Low information $\text{advantage}_{i(j)}$	-0.036	-0.014	-0.042	-0.012	-0.037	-0.008
	(0.097)	(0.092)	(0.097)	(0.092)	(0.109)	(0.104)
Controls	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year FE	Yes	Yes	Yes	Yes	No	No
Commuting zone × year FE	No	No	Yes	Yes	No	No
Industry × commuting zone × year FE	No	No	No	No	Yes	Yes
Observations	25,725	25,725	25,605	25,605	24,020	24,020
R^2	0.314	0.334	0.332	0.351	0.409	0.425

Notes: This table reports heterogeneous effects from equation 2 on asset growth. The sample covers the years 2010 to 2016. Post is equal to 1 in 2014 or after. $\text{Oil}_{i(j}$ is an indicator equal to 1 for entrepreneurs with spouses employed in the oil sector in both 2013 and 2014. The matched control group includes entrepreneurs with spouses employed in other sectors. Low $\text{information advantage}_{i(j)}$ is equal to 1 when the control entrepreneur within the matched pair has a top employee whose spouse works in the oil industry, giving the control firm indirect access to similar information and weakening the information gap relative to treated firms. The dependent variable is asset growth, defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. Controls include polynomials in firm age and in the entrepreneur's age, the entrepreneur's marital status, and the log number of owners. Standard errors are reported in parentheses and are clustered at the firm level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. The numbers of observations are rounded to the nearest five to comply with the confidentiality requirements of Statistics Canada.

B Computational Appendix

This appendix provides a complete description of the algorithms used to solve the model. The solution involves nested loops: an outer Brent iteration on the bond rate r (Algorithm 6), a direct wage computation from the no-arbitrage condition (no inner loop), value function iterations for entrepreneurs and workers (Algorithms 2–3), forward iteration for stationary distributions (Algorithm 4), and a top-level four-economy estimation (Algorithm 7). The Python implementation is contained in two files: `ge_solver.py` (core solver) and `run_all.py` (parameter interface and four-economy runner).

B.1 Grid Specification

The entrepreneur's state space (a, z) is discretized as follows. Wealth a is placed on a power-spaced grid of N_a points spanning $[a_{\min}, a_{\max}] = [0.001, 40.0]$:

$$a_i = a_{\min} + (a_{\max} - a_{\min}) \cdot \left(\frac{i}{N_a}\right)^{1.4}, \qquad i = 1, \ldots, N_a. \tag{26}$$

The power exponent of 1.4 provides denser coverage near low wealth levels where policy functions change rapidly. Productivity z follows a 5-state Tauchen approximation to the AR(1) process $\ln z_{t+1} = \varphi_z \ln z_t + \varepsilon_t$ with $\varepsilon_t \sim N(0, \sigma_z^2)$, using 3 standard deviations for the grid width. Because Jensen's inequality causes $[\exp(y)] > 1$ for a symmetric grid, the grid is *normalized* by dividing by its ergodic mean so that $\sum_i \pi_i z_i = 1$. This ensures that aggregate labor income equals the corporate wage bill and that average entrepreneurial TFP equals unity.

Capital k is discretized over N_k power-spaced points in $[0, k_{\max}]$, where $k_{\max}$ is the maximum feasible capital given the wealth grid upper bound and collateral constraints. Only values $k \leq a - c_{\min}$ are evaluated to satisfy the budget feasibility constraint.

For the worker problem, bond holdings b are placed on a power-spaced grid of $N_{a,w}$ points over $[0, 50]$. The worker state additionally includes productivity (5-state Tauchen with its own persistence and volatility parameters) and employment status $\varepsilon \in \{0, 1\}$.

B.2 Continuous Savings Optimization

A key computational innovation is the continuous optimization of entrepreneurial net savings s. For each state (a, z) and capital choice k, the feasible savings range is $s \in [-\bar{b}(k, w),\ a - k - c_{\min}]$, where $\bar{b}$ is the relevant collateral limit. Rather than evaluating a discrete grid of N_d savings points, we use Brent's bounded minimization algorithm—which combines golden-section search with parabolic interpolation—to find the optimal s to within an absolute tolerance of 10^{-6}.

For a given savings level s, the objective function is constructed as follows. Next-period

wealth depends on the realizations of productivity z' and (in E2) employment status ε:

$$a'(z', \varepsilon) = \tilde{w}(\varepsilon) + R^{\text{gross}}(k, z') + (1 + r)\, s. \tag{27}$$

The continuation value $V(a', z')$ is evaluated by linear interpolation on the wealth grid for each z', then aggregated across transitions in productivity z' and employment status ε under the Epstein-Zin risk aggregator:

$$\text{EV}^{1-\gamma_i} = \sum_{z'} P(z'|z) \sum_{\varepsilon} p(\varepsilon)\, \widetilde{V}(a'(z', \varepsilon),\, z')^{1-\gamma_i}. \tag{28}$$

The certainty equivalent then incorporates survival and entry:

$$\text{CE} = \left[\lambda \cdot \text{EV}^{1-\gamma_i} + (1 - \lambda) \cdot V^{\text{entry}}(a)^{1-\gamma_i}\right]^{1/(1-\gamma_i)}, \tag{29}$$

and the Bellman objective is $\text{val}(s) = [(1 - \beta)\, c^\rho + \beta\, \text{CE}^\rho]^{1/\rho}$ where $c = a - k - s$. The Brent algorithm typically converges in 10–15 function evaluations per (a, z, k) state, making it competitive with a grid search over $N_d = 20$–30 points while delivering a much smoother excess demand function $\text{ED}(r)$, which improves GE convergence.

B.3 Wage and Corporate Capital from No-Arbitrage

A simplification relative to many heterogeneous-agent models is that the wage is not determined by a separate clearing condition. Instead, for each candidate rate r, the no-arbitrage condition $r_K = r$ directly determines both the capital-labor ratio and the wage without iteration:

Algorithm 1: Wage and K_C from No-Arbitrage

Data: Rate r, TFP A, parameters

Result: Wage w, capital-labor ratio K_C/N_C, corporate capital K_C

1. $\text{KN} \leftarrow \left[(r + \delta)/(\alpha\, A)\right]^{1/(\alpha-1)}$ *// no-arbitrage pins K/N ratio*
2. $w \leftarrow (1 - \alpha)\, A\, \text{KN}^\alpha$ *// wage follows from CRS*
3. *After solving entrepreneur problem:*
   $N_E^{\text{lab}} \leftarrow$ aggregate entrepreneurial labor demand
   $N_C \leftarrow (1 - u)\, N_W - N_E^{\text{lab}}$
   $K_C \leftarrow \text{KN} \times N_C$
4. **return** w, KN, K_C

This eliminates the inner wage loop that would otherwise be needed, substantially reducing computation time. The entrepreneur VFI is solved once at $(r, w(r))$, and the resulting entrepreneurial labor demand determines N_C residually.

B.4 Entrepreneur Value Function Iteration

The entrepreneur's problem is solved by value function iteration over the state space (a, z) with a nested optimization over capital k (discrete grid) and net savings s (continuous via Brent's method). The algorithm is solved separately for each entrepreneur type $j \in \{L, H\}$.

Algorithm 2: Entrepreneur Value Function Iteration

Data: Bond rate r, wage w, TFP A, economy flag `has_wage` $\in \{\text{E1}, \text{E2}\}$, optional warm-start $V^{(0)}$
Result: Value functions $V_j(a, z)$, policies $k_j(a, z)$, $s_j(a, z)$, $c_j(a, z)$, for $j \in \{L, H\}$

1. **Discretize state space:**
   $\{a_i\}_{i=1}^{N_a} \leftarrow a_{\min} + (a_{\max} - a_{\min}) \cdot (i/N_a)^{1.4}$ on $[0.001, 40]$ *// power-spaced*
   $\{z_j, P_z\} \leftarrow \texttt{tauchen}(N_z, \varphi_z, \sigma_z)$; normalize so $\sum_j \pi_j z_j = 1$ *// Jensen correction*
   $\{k_m\}_{m=1}^{N_k} \leftarrow k_{\max} \cdot (m/N_k)^{1.5}$ where $k_{\max} = (a_{\max} + \theta_w w \cdot \mathbf{1}_{\text{E2}})/[1 - \theta_k(1 - \delta)]$
2. **Precompute:**
   $R_{m,j} \leftarrow \pi(k_m, z_j) + (1 - \delta)\, k_m$ for all (m, j) *// gross returns*
3. **if** `has_wage` **then**
   $\varepsilon_{\text{emp}} \leftarrow w - \tau_{\text{ui}}$, prob. $1 - u$; $\varepsilon_{\text{unemp}} \leftarrow \xi\, w$, prob. u; where $\tau_{\text{ui}} = u\xi w/(1 - u)$
   else $\varepsilon = 0$ with probability 1 *// E1: no spouse*
4. **foreach** type $j \in \{L, H\}$ with risk aversion γ_j **do**
   Initialize $V \leftarrow V^{(0)}$ (warm-start) or $V(a_i, z_j) \leftarrow 0.05\, a_i$
   while $t < I_{\max}^{\text{VFI}}$ **and not converged do**
   $V^{\text{entry}}(a_i) \leftarrow \left[\sum_z \pi_z\, V(\varphi \cdot a_i,\, z)^{1-\gamma_j}\right]^{1/(1-\gamma_j)}$ *// state-dependent*
   foreach state (a_i, z_j) **do**
   foreach capital level k_m with $k_m \le a_i - c_{\min}$ **do**
   $s_{\min} \leftarrow -[\theta_k(1 - \delta)\, k_m + \theta_w w \cdot \mathbf{1}_{\text{E2}}]$; $s_{\max} \leftarrow a_i - k_m - c_{\min}$
   $(s^*, c^*, \text{val}^*) \leftarrow \texttt{_solve_s_brent}(a_i, k_m, s_{\min}, s_{\max}, V, \ldots)$
   Objective for each s:
   $a'(z', \varepsilon) \leftarrow \tilde{w}(\varepsilon) + R_{m,z'}^{\text{gross}} + (1 + r)\, s$
   $\widetilde{V} \leftarrow \texttt{interp}(a', a\text{-grid}, V(\cdot, z'))$
   $\text{EV}^{1-\gamma_j} \leftarrow \sum_{z'} P(z'|z) \sum_{\varepsilon} p(\varepsilon)\, \widetilde{V}(a', z')^{1-\gamma_j}$
   $\text{CE} \leftarrow \left[\lambda \cdot \text{EV}^{1-\gamma_j} + (1 - \lambda)\big(V^{\text{entry}}(a_i)\big)^{1-\gamma_j}\right]^{1/(1-\gamma_j)}$
   $\text{val}(s) \leftarrow \left[(1 - \beta)\, c^{\rho} + \beta\, \text{CE}^{\rho}\right]^{1/\rho}$ where $c = a_i - k_m - s$
   Record $k_j(a_i, z_j) \leftarrow k^*$, $s_j(a_i, z_j) \leftarrow s^*$, $c_j(a_i, z_j) \leftarrow c^*$
   if $\|V^{(t+1)} - V^{(t)}\|_{\infty} < \text{TOL}_{\text{VFI}}$ **then break**
5. **return** $\{V_j, k_j, s_j, c_j\}_{j \in \{L,H\}}$, grids

B.5 Worker Value Function Iteration

Workers solve a standard Aiyagari-style consumption-savings problem with Epstein-Zin preferences and idiosyncratic unemployment risk. The state is (b, z_w) since employment is i.i.d.

Algorithm 3: Worker Value Function Iteration

Data: Bond rate r, wage w, TFP A, UI tax τ_{ui}

Result: Value function $V_w(b, z_w)$, savings policy $b'(b, z_w)$

1. $\{b_i\}_{i=1}^{N_{a,w}} \leftarrow b_{\min} + (b_{\max} - b_{\min}) \cdot (i/N_{a,w})^{1.4}$ on $[0, 50]$

2. $\{z_{w,j}, P_{z_w}\} \leftarrow \texttt{tauchen}(N_{z,w}, \varphi_w, \sigma_{e,w})$; normalize so $\sum_j \pi_j z_{w,j} = 1$

3. Income states:

$y_w(z_w, \varepsilon{=}1) = [w - \tau_{\text{ui}}]\, z_w$ *// employed, prob.* $1 - u$

$y_w(z_w, \varepsilon{=}0) = \xi\, w\, z_w$ *// unemployed, prob.* u

4. Initialize $V_w^{(0)}(b_i, z_{w,j}) \leftarrow [(1 - \beta)\, \bar{y}^{\,\rho}]^{1/\rho}$

5. while $t < I_{\max}^{\text{VFI}}$ **and not converged do**

foreach state $(b_i, z_{w,j})$ **do**

foreach savings choice $b' \in \{b_i\}$ with $b' \geq 0$ **do**

$\text{CE}_w(b') \leftarrow \left[\sum_{z'_w} P(z'_w|z_w) \sum_{\varepsilon'} p(\varepsilon')\, V_w^{(t)}(b', z'_w)^{1-\gamma_W}\right]^{1/(1-\gamma_W)}$

$b'^{*} \leftarrow \arg\max_{b'} \left[(1 - \beta)\, c_w^{\rho} + \beta\, \text{CE}_w(b')^{\rho}\right]^{1/\rho}$

if $\|V_w^{(t+1)} - V_w^{(t)}\|_\infty < \text{TOL}_{\text{VFI}}$ **then break**

6. return V_w, $b'(\cdot)$, grids

B.6 Stationary Distribution via Forward Iteration

The stationary distribution is computed by iterating the distribution forward using the policy functions and transition probabilities until convergence. This applies to both entrepreneurs and workers.

Algorithm 4: Stationary Distribution (Forward Iteration)

Data: Policy functions $k(\cdot), s(\cdot)$ or $b'(\cdot)$; grids; transition matrices P_z; shock distributions $(\varepsilon, p(\varepsilon))$

Result: Stationary distribution $\mu(a, z)$; aggregates K, N, B, D

1. Initialize $\mu^{(0)}(a_i, z_j) \leftarrow 1/(N_a \cdot N_z)$

2. while $\|\mu^{(t+1)} - \mu^{(t)}\|_\infty > 10^{-11}$ **do**

$\mu^{(t+1)} \leftarrow 0$

foreach state (a_i, z_j) with mass $\mu^{(t)}(a_i, z_j) > 0$ **do**

foreach (z', ε) realization **do**

$a' \leftarrow \tilde{w}(\varepsilon) + R^{\text{gross}}(k(a_i, z_j),\, z') + (1 + r)\, s(a_i, z_j)$

Find grid indices $\ell, \ell{+}1$ bracketing a'; compute weight ω

$\mu^{(t+1)}(\ell, z') \mathrel{+}= \lambda \cdot \text{mass} \cdot \text{prob} \cdot (1 - \omega)$; $\quad \mu^{(t+1)}(\ell{+}1, z') \mathrel{+}= \lambda \cdot \text{mass} \cdot \text{prob} \cdot \omega$

Turnover re-entry: $\forall\, (a_i, z_l)$, mass $(1 - \lambda)\, \mu^{(t)}(a_i, z_l)$ enters at $\max(\varphi \cdot a_i,\, \underline{a})$ with $z \sim \pi_z$

3. Normalize: $\mu \leftarrow \mu / \sum \mu$

4. Aggregate:

$K_j \leftarrow \sum_{a,z} \mu_j(a, z)\, k_j(a, z)$; $\quad N_j^{\text{lab}} \leftarrow \sum_{a,z} \mu_j(a, z)\, n^*(k_j, z; w)$

$B_j \leftarrow \sum_{a,z} \mu_j\, \max(-s_j, 0)$; $\quad D_j \leftarrow \sum_{a,z} \mu_j\, \max(s_j, 0)$

5. return $\mu, \{K_j, N_j, B_j, D_j\}$

B.7 Asset Market Excess Demand

Each evaluation of the excess demand $\text{ED}(r)$ requires: computing w from r via no-arbitrage, solving both entrepreneur types, the worker problem, and all stationary distributions.

Algorithm 5: Asset Market Excess Demand Evaluation

Data: Candidate rate r, parameters, TFP A, economy type
Result: Excess demand $\text{ED}(r)$, full equilibrium objects

Step 1: No-arbitrage pricing
$(w, \text{KN}) \leftarrow$ **Algorithm 1**(r, A) // $r_K = r$ *pins* w

Step 2: Entrepreneur problem
$\{V_j, k_j, s_j\}_j \leftarrow$ **Algorithm 2**(r, w, A); $\{\mu_j\}_j \leftarrow$ **Algorithm 4**
$N_E^{\text{lab}} \leftarrow N_E\,(f_L N_L + f_H N_H)$; $N_C \leftarrow (1-u)N_W - N_E^{\text{lab}}$; $K_C \leftarrow \text{KN} \times N_C$
$B_E \leftarrow N_E\,(f_L B_L + f_H B_H)$; $D_E \leftarrow N_E\,(f_L D_L + f_H D_H)$

Step 3: Worker problem
$(V_w, b'_w) \leftarrow$ **Algorithm 3**(r, w, τ_{ui})
$\mu_w \leftarrow$ `worker_stationary_dist`$(b'_w, \ldots)$
$S_w \leftarrow n_{\text{standalone}} \cdot \sum_{b,z_w} \mu_w(b, z_w)\, b'_w(b, z_w)$

Step 4: Asset market clearing
$\text{ED}(r) \leftarrow (K_C + B_E) - (S_w + D_E)$

Step 5: Walras' law check (diagnostic)
residual $\leftarrow (Y_C + Y_E) - (C_W + C_E + \delta\, K_C + \delta\, K_E)$

return $\text{ED}(r)$, all equilibrium objects

B.8 General Equilibrium Solver (Outer Loop)

The equilibrium rate r^* is found by Brent's method—which combines inverse quadratic interpolation with bisection fallback—applied to the excess demand function $\text{ED}(r)$. No inner wage loop is needed since w is pinned by no-arbitrage.

Algorithm 6: General Equilibrium: Brent's Method on r (Asset Market)

Data: Parameters, TFP A, economy type, optional r_{hint}
Result: Equilibrium rate r^* and all equilibrium objects

1. if r_{hint} provided **then** $[r_{\text{lo}}, r_{\text{hi}}] \leftarrow [r_{\text{hint}} - \Delta r,\ r_{\text{hint}} + \Delta r]$ with $\Delta r = 40\text{bp}$
else $[r_{\text{lo}}, r_{\text{hi}}] \leftarrow [0.001, 0.040]$

2. Evaluate endpoints:
$(\text{ED}_{\text{lo}}, \text{info}_{\text{lo}}) \leftarrow$ **Algorithm 5**(r_{lo}); $(\text{ED}_{\text{hi}}, \text{info}_{\text{hi}}) \leftarrow$ **Algorithm 5**(r_{hi})

3. **while** $\text{ED}_{\text{lo}} \cdot \text{ED}_{\text{hi}} > 0$ **do** *// no sign change*
Expand bracket in the direction of excess demand/supply (up to 8 attempts)
Re-evaluate the new endpoint via **Algorithm 5**
4. **Initialize Brent's method:**
$a \leftarrow r_{\text{lo}}, b \leftarrow r_{\text{hi}}$ with $|f(b)| \leq |f(a)|$; $c \leftarrow a$; `mflag` $\leftarrow$ `true`
Track best_info (smallest |ED| seen)
5. **for** $it = 1, \ldots, I_{\max}^{\text{GE}}$ **do**
if $|f(b)| < \text{TOL}_{\text{GE}}$ **then return** $r^* = b$ *// converged*
if IQI conditions satisfied **then** $s \leftarrow$ inverse quadratic interpolation step
else $s \leftarrow (a+b)/2$ *// bisection fallback*
$(\text{ED}_s, \text{info}_s) \leftarrow$ **Algorithm 5**(s)
Update best_info if $|\text{ED}_s| < |\text{best_ED}|$
Update bracket: $\{a, b\} \leftarrow$ standard Brent update ensuring $|f(b)| \leq |f(a)|$
if $|a - b| < 10^{-7}$ **then return** best_info *// bracket tight*
6. **return** best_info *// best solution found*

B.9 Four-Economy Estimation (Amplification)

The top-level estimation solves four general equilibrium problems sequentially, using warm-started interest rate brackets to reduce computation.

Algorithm 7: Four-Economy Comparison and Amplification

Data: Parameters, $A_{\text{base}} = 1.00$, $A_{\text{shock}} = 0.95$, bracket window $\Delta r = 40$bp
Result: Results for 4 economies, amplification ratio

Step 1: E1 Baseline (full bracket) *// no warm-start*
$\text{info}_{E1}^{\text{base}} \leftarrow$ **Algorithm 6**$(A = 1.00, \texttt{has_wage} = \text{False}, r_{\text{hint}} = \texttt{None})$; $r_1^* \leftarrow \text{info}_{E1}^{\text{base}}.\texttt{r}$
Step 2: E1 Shock *// warm-started from E1 base*
$\text{info}_{E1}^{\text{shock}} \leftarrow$ **Algorithm 6**$(A = 0.95, \texttt{has_wage} = \text{False}, r_{\text{hint}} = r_1^*)$
Step 3: E2 Baseline *// warm-started from E1 base*
$\text{info}_{E2}^{\text{base}} \leftarrow$ **Algorithm 6**$(A = 1.00, \texttt{has_wage} = \text{True}, r_{\text{hint}} = r_1^*)$; $r_3^* \leftarrow \text{info}_{E2}^{\text{base}}.\texttt{r}$
Step 4: E2 Shock *// warm-started from E2 base*
$\text{info}_{E2}^{\text{shock}} \leftarrow$ **Algorithm 6**$(A = 0.95, \texttt{has_wage} = \text{True}, r_{\text{hint}} = r_3^*)$
Step 5: Compute amplification
$\Delta Y_{E1} \leftarrow (Y_{E1}^{\text{shock}} - Y_{E1}^{\text{base}})/Y_{E1}^{\text{base}}$; $\Delta Y_{E2} \leftarrow (Y_{E2}^{\text{shock}} - Y_{E2}^{\text{base}})/Y_{E2}^{\text{base}}$
Amplification $\leftarrow \Delta Y_{E2}/\Delta Y_{E1}$ *// $> 1 \Rightarrow$ household economy amplifies shock*
return all results, amplification ratio

B.10 Algorithm Nesting Structure

Table B.5 summarizes the relationship between the algorithms and their convergence criteria.

B.11 Computational Parameters

B.12 Algorithm Correctness Validation

We verify the solution along three dimensions. First, *Walras' law*: at each r-evaluation, the goods market residual $Y - C - I$ is computed and confirmed to be small (typically < 0.01). Second, *no-arbitrage consistency*: the wage used in the entrepreneur VFI equals the wage implied by the no-arbitrage condition $r_K = r$ (exact by construction). Third, *distribution stationarity*: the forward-iteration distribution converges to a tolerance of 10^{-11}, and the distribution sums to unity after normalization.

Table B.5: Algorithm Nesting Structure

Layer	Algorithm	Method	Convergence
Top-level	Alg. 7: 4-economy comparison	Sequential warm-start	—
Outer	Alg. 6: Asset market (r)	Brent's method	$\lvert\text{ED}(r^*)\rvert < \text{TOL}_{\text{GE}}$
	$\hookrightarrow$ Alg. 5: ED evaluation	—	—
Inner	$\hookrightarrow$ Alg. 1: No-arb. pricing (w, K_C)	Direct formula	—
Core	$\hookrightarrow$ Alg. 2: Entrepreneur VFI	Value iteration	$\lVert V^{(t+1)} - V^{(t)}\rVert_\infty < \text{TOL}_{\text{VFI}}$
Core	$\hookrightarrow$ Alg. 3: Worker VFI	Value iteration	$\lVert V_w^{(t+1)} - V_w^{(t)}\rVert_\infty < \text{TOL}_{\text{VFI}}$
Dist.	$\hookrightarrow$ Alg. 4: Stationary dist.	Forward iteration	$\lVert \mu^{(t+1)} - \mu^{(t)}\rVert_\infty < 10^{-11}$

Table B.6: Computational Parameters (Precise Estimation)

Parameter	Symbol	Value	Description
Entrepreneur wealth grid	N_a	80	Power-spaced, $[0.001, 40]$
Capital grid	N_k	30	Power-spaced
Savings optimization		Continuous	Brent's method ($\lvert x_{\text{tol}}\rvert < 10^{-6}$)
Productivity states	N_z	5	Tauchen approximation
Worker wealth grid	$N_{a,w}$	150	Power-spaced, $[0, 50]$
Worker productivity states	$N_{z,w}$	5	Tauchen approximation
VFI tolerance	TOL_{VFI}	1×10^{-3}	Sup-norm on V
Max VFI iterations		800	
GE tolerance	TOL_{GE}	0.01	$\lvert\text{ED}(r^*)\rvert$
Max GE iterations		50	Brent's method on ED(r)
Interest rate bracket	$[r_{\min}, r_{\max}]$	$[0.001, 0.040]$	
Wage determination		No-arbitrage	$w(r)$: no inner loop needed
Distribution tolerance		10^{-11}	Sup-norm on μ