Entrepreneurs and (Risky) Spouses*

ISAAC HACAMO† and VALENTINA RUTIGLIANO§

November 25, 2026

ABSTRACT

How do spouses shape entrepreneurs' firm policies? A working spouse can insure the household against firm risk, but spousal income is itself risky and may fail precisely when it is most valuable: during recessions. Using an exogenous increase in spousal unemployment risk, we find that exposed entrepreneurs cut investment and leverage and extract liquidity through higher dividend payouts. The response is concentrated among riskier firms and entrepreneurs with the largest precautionary motives. A calibrated general-equilibrium model shows that countercyclical spousal unemployment risk can turn marriage from a stabilizer in normal times into an amplifier during downturns. Relative to a counterfactual all-singles economy, an economy in which entrepreneurs are married to workers produces higher output in good times but experiences larger output losses following a negative productivity shock.

JEL classification: D13; G51; G32; L26

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

*We thank seminar participants at Deakin University, Federal Reserve Bank of Chicago, Indiana University, Rice University, Southern Methodist University, Texas Tech University, University of Delaware, University of Massachusetts at Amherst, UNSW Sydney, and University of Queensland; and conference participants at Chicago Entrepreneurial Finance Workshop, 40th Mitsui Finance Symposium: Private Capital Markets, Nordic Household Finance Summit, NFA, RAPS/RCFS Europe Conference 2026, UGA Fall Finance Conference, and UNSW Corporate Finance Workshop. We also thank Gene Amromin, Daniel Carvalho, Mayur Choudhary, Paul Freed, Xiaoyang Li, Mariacristina De Nardi, Kasper Meisner Nielsen, Alessio Piccolo, Gordon Phillips, Eddie Shore, Wing Wah Tham, and Daisy Wang for very helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of Statistics Canada.

†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 U.S. entrepreneurs are married, yet entrepreneurs are often studied as standalone decision-makers.[1] Ignoring spouses overlooks a critical economic mechanism: household insurance. Because entrepreneurial income and spousal labor income are pooled in a common household budget constraint, spousal income can cushion household consumption when the business performs poorly. But this insurance is itself risky. A spouse is more likely to lose a job in a recession, precisely when the firm is also likely to be distressed. Spousal insurance may therefore fail in the states in which coverage is most valuable. If this insurance fails in bad aggregate states, marriage can amplify macroeconomic shocks rather than dampen them. Despite the importance of this insurance mechanism, our understanding of how spouses shape entrepreneurial decisions remains limited.

We first provide a reduced-form test of whether entrepreneurs' firm policies respond to the insurance value of spousal income. We identify this channel using an exogenous increase in spousal unemployment risk, which reduces the strength of spousal insurance. Exposed entrepreneurs de-risk. They cut investment and leverage. These adjustments occur even when the spouse remains employed, indicating that entrepreneurs respond to heightened spousal unemployment risk. Exposed entrepreneurs also increase dividend payouts, consistent with precautionary liquidity extraction rather than tighter borrowing constraints.[2] Moreover, the effects are stronger among entrepreneurs with the largest precautionary savings motives—entrepreneurs with younger children or older ones—and among firms with more volatile cash flows. This evidence is consistent with de-risking in response to less reliable spousal insurance.

We then build and calibrate a heterogeneous-agent general-equilibrium model with married households and endogenous occupational choice. Spousal income insures household consumption and relaxes collateral constraints, so in normal times married entrepreneurs operate larger, more leveraged firms than comparable singles. But unemployment risk is countercyclical: recessions weaken spousal insurance precisely when firm cash flows are weakest. Anticipating this, married entrepreneurs de-risk

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

[2] If borrowing constraints were the primary mechanism, firms would instead retain cash.

in recessions before any job loss occurs, amplifying the output loss from a negative productivity shock relative to an all-singles economy.

To provide reduced-form evidence that entrepreneurial firm policy responds to the insurance value of spousal labor income, we exploit the abrupt collapse in global oil prices in 2014. Oil prices fell sharply in 2014, after global supply began to exceed demand, driven primarily by rapid growth in U.S. shale production.[3] The decline accelerated when OPEC chose not to cut production in late 2014, reportedly to preserve market share relative to non-OPEC producers.[4] Global oil prices fell from $105 to $50 per barrel in the second half of 2014. The collapse, which was largely unexpected by Canadian households and firms, and triggered a recession and widespread unemployment in oil-producing regions during 2015–2016. We use the oil price collapse to design an empirical design in which some entrepreneurs experience a plausibly exogenous increase in spousal unemployment risk. The setting also offers a reliable placebo test that further strengthens our identification strategy.

We define treated entrepreneurs as those whose spouses were employed in the oil sector in 2013–2014, before the oil price plunge, and were unexpectedly exposed to a sharp increase in unemployment risk. For each treated entrepreneur, we construct a control by matching to an entrepreneur in the same four-digit industry, gender, and narrow geographic area, whose spouse had similar pre-shock income but worked outside the oil sector. This design compares firms operating in the same narrowly defined product and local labor markets and facing similar pre-shock household budget constraints, while differing in their exposure to the rise in spousal unemployment risk. Although we do not match on pre-shock spousal unemployment risk, spousal layoff probability is balanced between treated and control entrepreneurs, indicating that the two groups had similar pre-shock levels of spousal insurance. Further, treated and control entrepreneurs are well balanced across a broad set of pre-shock characteristics, including the entrepreneur's income, firm assets, revenues, leverage, and firm age.

We also implement a placebo test based on entrepreneurs whose spouses had

[3] See Krauss (2014), who reports that the boom in American energy production was keeping the market "flush with crude" and putting downward pressure on oil prices.

[4] Jahn (2014) reports that OPEC ministers viewed a production cut as likely to reduce OPEC's market share while allowing non-OPEC producers to make up the difference.

been employed in the oil sector but *luckily* exited in the months immediately before the shock. This group provides a natural benchmark: these entrepreneurs had spouses with similar prior exposure to, and selection into, the oil industry, but their spouses were not subsequently exposed to the increase in unemployment risk. Thus, if endogenous selection into the oil industry shaped entrepreneurial outcomes, this placebo group should exhibit similar effects to treated ones.

Studying how spousal unemployment risk affects entrepreneurial firms requires unusually rich data. One must observe the entrepreneur's household, the spouse's labor market exposure, and the financial decisions of the entrepreneur's firm. We overcome these challenges using 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 spouses' employment histories to detailed financial information on the entrepreneur's firm.[5] We start by showing that the central empirical pattern underlying our hypothesis holds for the universe of Canadian entrepreneurs: increases in the layoff rate of the spouse's industry are associated with lower asset growth in the entrepreneur's firm. This evidence not only motivates our causal analysis, but also supports the external validity of our causal setting, which focuses on a single industry.

We then turn to our main results. We first show that Canadian workers exposed to the 2014 oil shock were more likely to experience job loss and earnings losses. 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 its well-established impact on aggregate unemployment documented in the literature. An event-study analysis shows no differential pre-trends for treated and control spouses, further strengthening the validity of our setting.

We then examine the firm decisions of entrepreneurs married to these workers,

[5]The Canadian setting is also particularly well suited to our question as the vast majority of entrepreneurs are part of households. Figure 1 shows that, by age 35, more than 80 percent of entrepreneurs in Canada are married or cohabiting.

relative to control entrepreneurs. We first document that exposed entrepreneurs reduce firm size. Asset growth declines by over 4 percentage points, a decline equal to about 8 percent of the standard deviation. This reduction is accompanied by lower risk exposure. Leverage declines by almost 3 percentage points. When we restrict the sample to spouses who remained employed, the point estimates are almost identical. Additionally, *lucky* workers who moved from the oil sector to a non-oil sector in the months before the shock are no more likely to experience unemployment, and their entrepreneur spouses' firm outcomes are indistinguishable from those of the control group. These results show that spousal unemployment risk shapes entrepreneurial firm policy and are consistent with both precautionary responses to weaker household consumption insurance and tighter borrowing constraints.

However, the two mechanisms have opposite implications for dividend payouts. Tighter borrowing constraints predict lower payouts, whereas precautionary responses to weaker household insurance predict higher payouts. Since our data measure dividend payouts even for small firms, we can use them to separate these two mechanisms. We find that entrepreneurs with spouses in the oil industry are 3 percentage points more likely to increase dividend payouts. Total compensation rises by 5 percent and is entirely driven through dividends as salaries do not change. This effect is economically large and supports the role of precautionary responses to weakened household insurance in explaining our findings.

We further investigate the precautionary mechanism using the rich demographic data in our dataset. Older entrepreneurs are 6 percentage points more likely to reduce asset growth than younger entrepreneurs, and older female entrepreneurs are a further 8 percentage points more likely to do so. The response also increases with the number of young children: households with two young children respond more than households with one, which respond more than households with none. These are the groups with the strongest savings motives, supporting our interpretation that our fidings are driven by a precautionary mechanism.

The deterioration of spousal income insurance should also lead to stronger precautionary responses when firm cash flows are riskier: more volatile cash flows expose entrepreneurs to greater downside risk, leading them to cut investment more aggres-

sively following an increase in spousal unemployment 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 large declines in investment following an increase in spousal unemployment risk. We also show that the negative effect on asset growth is substantially larger among young firms.[6]

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.[7] We also observe small and non-significant effects on sales growth, further weakening the possibility that our effects are driven by local demand effects.

Second, we rule out the possibility that business relationships with the oil industry explain our findings. Managers in the oil sector are at least as exposed to the oil shock as other employees, yet entrepreneurs married to managers experience slightly smaller declines in asset growth following the shock. If business opportunities were transmitted through the spouse's network, entrepreneurs married to managers would experience larger declines in firm growth. Additionally, the null effects for entrepreneurs married to *lucky* spouses are difficult to reconcile with an explanation based on persistent business relationships with the oil sector. It is unlikely that exposure would suddenly disappear after a spouse's industry switch.

Third, we assess whether informational spillovers within the household can account for our results. Three patterns are difficult to square with this explanation. The effects persist through 2016, well after the 2014 oil price collapse had become widely known in Canada. Moreover, 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: effects are larger in oil-producing

[6]Defined as firms that are at most five years old in 2014.

[7]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.

provinces. Finally, we redefine the control group to include only entrepreneurial firms that employ a top non-owner manager whose spouse works in the oil industry. 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 no evidence of such attenuation.

We then develop a heterogeneous-agent general-equilibrium model in which households choose between corporate employment and entrepreneurship. Entrepreneurs, who are heterogeneous in productivity, finance firms with their own wealth and external borrowing, subject to a collateral constraint. In couples in which one spouse runs a firm and the other works, the worker spouse's income plays a dual role: it expands pledgeable collateral, relaxing borrowing constraints, and supplements household resources, providing insurance that single entrepreneurs lack.

The model also features countercyclical unemployment risk: recessions simultaneously reduce TFP and increase the probability that the worker spouse loses their job, weakening the collateral and insurance channel that marriage provides. This tension between insurance and risk exposure generates a sharp precautionary response. Married entrepreneurs scale back capital and leverage even when their spouse remains employed, anticipating the possibility of future insurance and collateral loss. Quantitatively, the model produces a 4.7% output amplification in the couples economy relative to a counterfactual all-singles benchmark. The model-implied difference-in-differences estimates are consistent with our reduced-form results.

Two additional results help us understand how married entrepreneurs amplify negative productivity shocks. First, lowering the fraction of pledgeable spousal income that entrepreneurs can use to relax borrowing constraints has a negligible effect (i) the capital and leverage response, and (ii) on the output amplification of negative shocks. Consistent with our reduced-form results, precautionary responses to the weakening of spousal insurance are the main driver of the model results. Second, when we shut down cyclical unemployment risk, married entrepreneurs increase capital and leverage in response to negative productivity shocks relative to their single counterparts. This reversal again highlights the importance of spousal insurance in explaining the amplification of negative shocks during downturns.

Together, our evidence shows that spousal income provides valuable household insurance in normal times, but this insurance becomes fragile in downturns. When spousal unemployment risk rises, entrepreneurs reduce firm scale and risk exposure, and these private adjustments amplify aggregate shocks. The results highlight a household channel through which labor-market risk shapes firm policy and macroeconomic dynamics.

Related literature. Our work relates to three strands of literature. First, we contribute to the literature on entrepreneurial risk-taking. Entrepreneurial households bear substantial idiosyncratic risk (Moskowitz and Vissing-Jørgensen, 2002). Entrepreneurs' lack of diversification affects real investment (Panousi and Papanikolaou, 2012), capital structure (Chen, Miao, and Wang, 2010), and risk-taking (Chen and Strebulaev, 2019). Providing protection against adverse outcomes can make this risk easier to bear. Providing unemployment insurance to unemployed entrepreneurs in case of failure (Hombert et al., 2020) and extending job-protected leave (Gottlieb, Townsend, and Xu, 2022) encourage entrepreneurial experimentation. Policies that improved coverage of idiosyncratic risk impacts not only inequality dynamics within and across generations but also macro capital accumulation (Silva and Townsend, 2025).[8] We identify spousal labor income as a source of private insurance for entrepreneurs. Spousal insurance supports entrepreneurial risk-taking in good times, yet this insurance is itself exposed to cyclical fluctuations. Whereas prior work emphasizes public insurance and wealth as buffers against entrepreneurial risk, we show that insurance provided within the household is state dependent and can transmit labor-market downturns to entrepreneurial firms.

Second, our paper contributes to the literature on household self-insurance against income and unemployment risk. Under incomplete markets, households smooth consumption by accumulating precautionary savings and shifting their financial portfolios

[8] Both current and expected wealth shocks also spur firm creation, consistent with increased willingness to take risk. Hurst and Lusardi (2004) find that entry increases in response to both past and future inheritances, while Bianchi and Bobba (2013) show that expected future cash transfers increase entry more than transfers received today. A related literature shows that receiving inherentences increases individuals likelihood of becoming entrepreneurs (Holtz-Eakin, Joulfaian, and Rosen, 1993; Hurst and Lusardi, 2004; Andersen and Nielsen, 2012).

toward safer or more liquid assets (Carroll, Dynan, and Krane, 2003; Bayer et al., 2019; Fagereng, Guiso, and Pistaferri, 2018; Catherine, 2022; Juelsrud and Wold, 2025). Couples can further share risk by adjusting labor supply when one spouse's income falls (Hess, 2004; Blundell, Pistaferri, and Saporta-Eksten, 2016; Mankart and Oikonomou, 2017; De Nardi, Fella, and Paz-Pardo, 2025). The canonical example is the *added-worker* effect, whereby the secondary earner (typically the wife) enters the labor market following the primary earner'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 entrepreneurs responds to household income risk by altering firm policies. Firm adjustments transmit risk-management motives beyond the household, amplifying the propagation of macroeconomic shocks.

Finally, we contribute to a large literature on the propagation of economic shocks 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). In these studies, shocks travel along formal economic relationships among firms and financial intermediaries. We document a transmission channel that operates outside any such relationship: the household. This household-to-firm channel is a previously unmeasured node through which macroeconomic shocks can spread.

2 Institutional Setting

2.1 The 2014 oil price shock in Canada

The oil and gas sector is a cornerstone of the Canadian economy, accounting for 7.5% of GDP in 2013 (Fremeth, Holburn, and Loudermilk, 2016). Its economic importance varies substantially across provinces. Alberta is the center of Canada's oil industry and accounted for more than 75% of national crude oil production in 2013. Saskatchewan, Newfoundland and Labrador, and British Columbia are also closely exposed through

production and links to the broader oil and gas sector. We refer to these four provinces as the oil-exposed regions.[9]

The episode we study is the 2014 collapse of global oil prices. After exceeding USD 100 per barrel in June 2014, WTI crude fell to below USD 50 by early 2015 as expanding U.S. shale supply outpaced global demand. The shock brought the preceding resource boom to an abrupt end. Canadian oil producers scaled back investment, reducing labor demand both in extraction and in industries serving the sector.

Figure 2 documents the resulting reversal in economic activity. Panel A plots real GDP per capita. Between 2010 and 2014, GDP per capita increased by approximately 7% nationwide and 11% in oil-exposed regions. This reflects a period of high investment in oil sands and offshore production, alongside strong demand for related industrial support. After mid-2014, the pattern reversed: GDP per capita in oil-exposed regions declined by roughly 1.5% and 6% nationwide and in oil regions, respectively, between the end of 2014 and mid-2016. Panel B shows the corresponding evolution of unemployment. Before the shock, unemployment in oil-exposed provinces remained below 6%, more than one percentage point below the national rate. It rose rapidly after mid-2014 and eventually exceeded the national figure.

Figure 3 presents the evolution of unemployment across Canadian Economic Regions (which are sub-provincial divisions). Panel A shows that in 2013, the regions closely related to oil activity were among the tightest labor markets in Canada. Unemployment was below 5% in Alberta and was also relatively low in parts of Saskatchewan and northeast British Columbia connected to oil. Newfoundland and Labrador had a higher unemployment rate, although the region most closely connected to offshore production performed better than the rest of the province. By 2016 (Panel B), this ordering had reversed. The northern Alberta oil-sands region and Calgary moved from unemployment rates below 5% to rates above 9%, with similar, though less pronounced, deteriorations in oil-connected areas of Saskatchewan, British Columbia, and New-

[9]Production in Alberta is dominated by oil sands, while Saskatchewan produces primarily from conventional oil fields and Newfoundland and Labrador produces offshore. British Columbia's oil and gas fields are concentrated in the northeast, which is closely integrated with the broader Western Canadian energy sector. Other regions of Canada are indirectly exposed through industries that supply the oil sector, including construction, engineering, transportation, equipment, and specialized industrial services.

foundland and Labrador. Thus, the regions that had benefited most from the resource boom experienced the largest reversals after oil prices collapsed.

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.[10]

Canada's unemployment insurance system provides partial income support following separation. 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).

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

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 entrepreneurial firms respond to spousal unemployment risk requires a source of variation that is external to the firm and affects the entrepreneur's household without directly influencing firm fundamentals. This is challenging because, in many entrepreneurial households, unemployment 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 unemployment risk for workers in the oil sector. For our purposes, it is crucial that the 2014 oil-price plunge was 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). Because the collapse was largely unexpected, it provides a credible source of variation in unemployment risk for our identification strategy.

4.1 Construction of treatment and control groups

We aim to identify whether entrepreneurial firms respond to shocks to the insurance value of spousal labor income. This requires variation in household income risk that is external to the firm and does not directly affect firm fundamentals. This is challenging because, in many entrepreneurial households, labor income risk is tightly linked to the firm itself. We therefore focus on entrepreneurs whose spouses are employed outside the firm and are not involved in its operations. This allows us to study how firm decisions respond to changes in the household's broader financial environment, rather than to shocks to the firm's own productivity, demand, or financing conditions.

We exploit such variation from the 2014 oil price collapse, which generated a sharp and largely unexpected increase in unemployment risk for workers in the oil sector. Because the collapse was largely unanticipated, it provides a credible source of variation in spousal unemployment risk for our identification strategy. The key comparison is between entrepreneurs whose spouses were exposed to this increase in unemployment risk and observationally similar entrepreneurs whose spouses worked outside the oil sector.

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 before the shock and were exposed to its consequences at the onset of

the collapse. 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.[11]

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

We implement one-to-one matching to construct a comparable control group. Matching is performed in 2013, before 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, or "caliper." We adopt a deliberately parsimonious matching strategy based on four pre-shock characteristics: we match exactly on four-digit industry, entrepreneur sex, and commuting zone,[12] and impose a 25-log-point caliper on spousal earnings.

This matching strategy is designed to compare entrepreneurs operating in the same narrowly defined product and local labor markets and facing similar pre-shock household budget constraints. Matching on the entrepreneur's four-digit industry holds fixed the firm's sectoral exposure. Matching on commuting zone 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 changes in local income, labor market conditions, or demand for goods and services. By comparing treated and control entrepreneurs

[11] In robustness checks, we verify that the results are not driven by differential divorce rates after the shock by restricting the sample to couples who remain married.

[12] 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.

within the same commuting zone, we hold fixed these local recessionary conditions while isolating differences in exposure to spousal unemployment risk.

Importantly, we do not match directly on pre-shock spousal unemployment risk. This allows us to test whether the oil shock changed the reliability of spousal insurance rather than mechanically imposing balance on the outcome most closely related to treatment. In the data, however, treated and control spouses have similar pre-shock layoff probabilities, indicating that the two groups had similar levels of spousal insurance before the shock. Treated and control entrepreneurs are also well balanced across a broad set of pre-shock characteristics, as discussed below.

4.2 Sampling

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.

4.3 Balancing tests

Table 1 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 4 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.

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.

4.4 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:

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

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 spousal unemployment 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.

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, relative to control spouses. Figure 5 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 treated and control groups. The pattern changes starting in 2014. Layoff probability increases by about 2.3 percentage points in that year and continues to rise over 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. The two indicators capture different phases of job displacement: layoffs measure employer-reported involuntary separations, while unemployment insurance receipt captures periods spent without work. Not all separations result in unemployment, since 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 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 suggests that early separations were often absorbed by the local labor market, with unemployment rising only as conditions deteriorated further.

Table 2 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 show that spouses employed in the oil sector experienced a substantial increase in job loss and unemployment risk, along with a decline in earnings, following the 2014 oil price crash. Treated households therefore faced a meaningful deterioration in spousal income insurance, setting the stage for examining how entrepreneurs adjust their firms' decisions.

5.2 Do entrepreneurs internalize household shocks?

We now turn to the paper's central question: do entrepreneurs internalize the weakening of spousal income insurance when making business decisions? If entrepreneurs treat their firms as separate from the household, spousal unemployment risk should not affect firm policies. But if household resources are pooled, greater uncertainty about the spouse's labor income may lead entrepreneurs to adjust investment, payout, and capital structure decisions to smooth household consumption.

Figure 6 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 curtailing firm

scale.

Panels (c) and (d) 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 to 2 percentage points in 2016. The direction of the effect is informative because leverage could, in principle, move either way. If debt declines less than assets, leverage could increase mechanically. In contrast, if households have a strong incentive to lower the firm's risk exposure, they could cut debt by more than assets, leading to lower leverage. A decline in debt could also reflect reduced credit supply if spousal income is used as collateral to relax borrowing constraints. The observed decline in leverage is therefore consistent with both active de-risking by entrepreneurs and a tightening of borrowing capacity.

Panels (e) 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 (f), 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 a precautionary response at the household level.

Table 3 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. Finally, 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. The most demanding specification, reported in column (6), implies a decline of 10 percent relative to the mean.

Table A.3 complements the extensive-margin evidence on dividend payments by examining the amount of resources entrepreneurs withdraw from the firm. Columns (1)–(4) presents results for total compensation, defined as the sum of dividends and salary paid to the entrepreneur, scaled by firm sales. Compensation rises by roughly 5 percentage points following the shock. Columns (5)–(8) and (9)–(12) decompose this increase into salary and dividend payments, respectively. The increase is concentrated in dividends, which rise by 3.9–4.3 percentage points. Salary payments increase by only 0.7–0.9 percentage points and are not statistically significant in the specifications with the richest fixed effects. This evidence indicates that treated entrepreneurs increase the amount of resources withdrawn from the firm, consistent with using the firm as a source of liquidity, rather than merely shifting compensation from salary to dividends.

Table A.4 turns to firm operating outcomes. Columns (1)-(4) show that the effect on sales growth is small and imprecisely estimated. This makes a direct demand shock to the entrepreneur's business an unlikely explanation for the financial adjustments documented above. Profitability, however, declines: profit margins fall by 2.4-2.8 percentage points following the shock. Figure A.2 shows that the decline in profitability emerges gradually after the shock, consistent with reduced investment eroding firm performance over time rather than with an immediate adverse shock to the business. The contraction in assets does not appear to represent a shift toward a smaller but more efficient firm. The deterioration in profitability also provides little support for the view that entrepreneurs reduce investment because they have superior information about the downturn. We return to these alternative interpretations in Section 5.3.

Table A.5 provides further evidence that the investment response is not concentrated among marginal businesses. Outcomes are similar between firms below and above

the median of the pre-shock size distribution, with slightly larger and more precise estimates among larger firms. Among firms with above-median pre-shock profitability, asset growth declines by 7.0 percentage points in our most demanding specification. By contrast, the corresponding estimates for below-median-profitability firms is 2.4 percentage points and is imprecisely estimated. The contraction is therefore, if anything, concentrated among firms that were performing relatively well before the shock.

We next examine whether the investment response varies with the entrepreneur's role in household income. We define the entrepreneur's income share as the entrepreneur's pre-shock income divided by total household income, and construct terciles of this variable separately within the treatment and control groups in Table 4. The decline in asset growth is concentrated in balanced-income households, defined as households in the middle tercile of the entrepreneur's income share. In this group, treated firms reduce asset growth by 6.3 to 7.7 percentage points after the shock. By contrast, the effects are smaller and statistically insignificant when the entrepreneur contributes a small or very large share of household income. This pattern is consistent with the insurance mechanism. When the entrepreneur's income share is very high, spousal insurance is limited, making these households closer to single-entrepreneur households. When the spouse's income share is very high, the entrepreneur contributes less to household consumption, reducing the need to adjust firm policy. The largest effects should therefore arise when both spouses contribute meaningfully to household income, as we observe in the data.

Risk without realized income shocks—. The distinction between increased income risk and realized income shocks is central to interpreting our findings. If entrepreneurs adjust firm policies only after their spouse's income falls, the results could reflect the direct effect of lower household income. We re-estimate the main specifications after excluding spouses in the bottom quartile of earnings growth over the shock period. This restriction removes households with the largest negative earnings realizations, including but not limited to cases of job displacement. In the remaining sample, spousal earnings growth is positive for both treated and control households: median earnings growth is 2.4 percent for controls and 2.5 percent for treated spouses. Figure A.1 shows

that among these households, entrepreneurs exposed to the oil sector through their spouse reduce asset growth, increase dividend payouts, and lower leverage after the shock, with estimates close to the baseline.

5.3 Mechanisms

Precautionary response—. 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 and the household as jointly connected. When spousal income becomes less reliable, the household has less insurance against poor firm outcomes. Entrepreneurs may therefore respond by reducing the firm's risk exposure and extracting liquidity from the firm.

The increase in dividend disbursements documented in the previous section is consistent with this mechanism. Dividend payouts allow entrepreneurs to move resources out of the firm and into the household when the insurance value of spousal income weakens. Because this pattern arises even when spouses do not lose their jobs or experience income losses, it is difficult to explain with tighter borrowing constraints alone. Tighter borrowing constraints would instead predict that firms retain internal funds and reduce dividend payouts.

This mechanism implies that the response should be stronger among households with greater precautionary motives. When the household places greater weight on protecting consumption against adverse income realizations, a deterioration in the reliability of spousal income should induce a larger reduction in firm investment. Table 5 examines this prediction using demographic characteristics associated with households' exposure to downside income risk. Panel A examines heterogeneity by age and gender. To compare spouses at similar stages of the life cycle, the panel restricts the sample to couples whose ages differ by no more than 20 years. Columns (1) and (2) identify older couples as those in which both the entrepreneur and the spouse are above the 75th percentile of the age distributions. Among these couples, asset growth declines by an additional 6.6–6.7 percentage points relative to other couples. 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. Asset growth among these entrepreneurs declines by an additional 9.1–9.4 percentage points, consistent with evidence of gender differences in tolerance for financial risk (Charness and Gneezy, 2012).

The investment response is also stronger in households with young children. Children increase Households with young children should therefore place greater value on liquidity when spousal income becomes less secure. Panel B of Table 5 splits households by the presence of children younger than eight in 2014. Asset growth declines by 4.0–4.8 percentage points among entrepreneurs with no young children, by 6.8–7.2 percentage points among those with at least one young child, and by 10.6–11.6 percentage points among those with at least two young children. These results suggest that weaker spousal insurance has stronger effects when households face larger fixed commitments and less flexibility to cut down on consumption.

A second implication of the insurance mechanism is that household should adjust more strongly when firm cash flows are riskier. When firm cash flows are volatile, weaker spousal insurance leaves the household more exposed to downside firm outcomes. Entrepreneurs should therefore have a stronger incentive to reduce investment and risk exposure. By contrast, when firm cash flows are stable, the loss of spousal insurance is less likely to require large adjustments in firm policy.

Table 6 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, reported in 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, where asset growth declines by about 5 to 7 percentage points following the increase in spousal unemployment risk.

We also consider firm age as an additional proxy for firm-level risk. Younger firms

tend to have more volatile cash flows than established firms, making the loss of spousal insurance more consequential for household consumption risk. Table A.8 shows that the decline in asset growth following the shock is substantially larger among young firms, defined as firms that are at most five years old in 2014. 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, compared with 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.9 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 × 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 8 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.[13] 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 7 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

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

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—. An alternative interpretation is that treated entrepreneurs receive earlier or more precise information about the impending economic crisis, consistent with information spillovers between spouses (Delavande, Koşar, and Zafar, 2025). Under this interpretation, entrepreneurs reduce investment because they anticipate worsening economic conditions rather than because the shock raises household labor-income risk.

The information channel yields several distinct implications. First, an informational advantage should matter most when information about the shock is scarce. This prediction is difficult to reconcile with the timing and geography of firms' responses: the effects persist through 2016, well after the collapse in oil prices had become widely known and incorporated into public expectations; in addition, the effects are larger in oil-producing provinces, where oil-sector developments are especially salient and closely monitored.

Second, the information channel predicts a stronger response when the spouse is particularly well positioned to acquire and transmit relevant information. This interpretation is further weakened by the evidence on manager spouses presented in Table 8 and discussed above: entrepreneurs married to managers, who should be best positioned to transmit information, do not exhibit stronger investment responses.

Conversely, the treatment effect should attenuate when control entrepreneurs have access to an alternative conduit of oil-sector information. Table A.10 identifies control firms whose highest-ranked non-owner employees have a spouse working in the oil sector. These employees provide a plausible channel through which information about developments in oil could reach the entrepreneur even when the entrepreneur's own

spouse does not work in the sector. The treatment effect remains similar in this sample, providing little evidence that equalizing access to oil-sector information narrows the treated-control gap.

Finally, if treated entrepreneurs reduce investment because superior information allows them to anticipate the downturn and adjust more effectively, this advantage should eventually be reflected in relative firm performance. Instead, Figure A.2 shows that profitability deteriorates progressively following the shock. The investment response therefore does not appear to represent an effective anticipatory adjustment based on superior information.

Divorce—. Divorce is another potential channel through which an increase in spousal labor income risk could affect entrepreneurial firms. If the oil shock raises the probability of marital dissolution, the estimated firm responses could reflect the financial consequences of divorce rather than precautionary adjustments within continuing households. Our baseline specifications control for contemporaneous marital status, accounting for differences in firm outcomes associated with marital status. However, because marital dissolution may itself be caused by the spouse's employment shock, we also re-estimate our main specifications in Table A.6 on the subsample of couples who remain married throughout the sample period. The estimated effects on asset growth are essentially unchanged, indicating that the firm-level responses we document are not driven by marital dissolution.

6 Macro Model

Spousal insurance has opposing implications for the propagation of macro shocks. By supporting household consumption and relaxing financial constraints, spousal income can stabilize entrepreneurial firms in normal times. But when spousal unemployment risk rises in downturns, this support weakens precisely when firms are most vulnerable, inducing entrepreneurs to de-risk and potentially amplifying aggregate shocks. We therefore introduce a calibrated heterogeneous-agent general-equilibrium model to quantify these forces relative to a conventional benchmark in which all entrepreneurs are single.

The framework follows Cagetti and De Nardi (2006), Quadrini (2000), and Moll (2014). More concretely, the model features single and married households whose members choose each period between paid employment and entrepreneurship. Entrepreneurs differ in productivity and finance decreasing-returns firms with their own wealth and external borrowing, subject to collateral constraints. For married entrepreneurs, spousal income increases pledgeable household income and insures household consumption when firm income is low. Recessions lower firm productivity and raise the probability that the worker spouse loses their job, allowing the model to capture how the same household structure that supports entrepreneurship in normal times can generate de-risking and amplify downturns.

6.1 Model

6.1.1 Environment

Time is discrete and infinite. The economy contains a unit mass of agents. A fraction m live in married couples (each couple counted as one unit of two individuals) and $1-m$ are single. All agents share Epstein–Zin preferences with risk aversion γ, IES ψ, and discount factor β:

$$V_t = \left[(1-\beta)\, c_t^{\rho} + \beta\, \mathrm{CE}_t^{\rho}\right]^{1/\rho}, \qquad \mathrm{CE}_t = \left[E_t\left(V_{t+1}^{1-\gamma}\right)\right]^{1/(1-\gamma)}, \tag{3}$$

where $\rho = 1 - 1/\psi$. Couples maximize a single household value function over joint consumption and occupational choices; there is no intra-household bargaining.

6.1.2 Productivities and Matching

Each agent carries two productivity draws. A *permanent worker type* $z_w \in \{z_w^L, z_w^H\}$ determines corporate-sector earnings $w \cdot z_w$. A separate *entrepreneurial productivity* z_e follows a persistent AR(1) process in logs,

$$\ln z_{e,t+1} = \rho_z \ln z_{e,t} + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_z^2), \tag{4}$$

discretised into N_Z states via Tauchen's method. Because z_w and z_e are independent, high-wage workers select into entrepreneurship only when their entrepreneurial draw is sufficiently favorable.

Couples are formed with assortative matching on z_w. Let $p = \Pr(z_w = z_w^H)$. The joint distribution of spousal types is parameterised by $\rho_m \in [0, 1]$:

$$\Pr(H, H) = p^2 + \rho_m\, p(1-p), \tag{5}$$

$$\Pr(L, L) = (1-p)^2 + \rho_m\, p(1-p), \tag{6}$$

$$\Pr(H, L) = \Pr(L, H) = p(1-p)(1-\rho_m). \tag{7}$$

Random matching obtains when $\rho_m = 0$; perfect positive assortative matching when $\rho_m = 1$.

6.1.3 Production

A competitive corporate sector operates CRS technology $Y_C = A K_C^{\alpha} N_C^{1-\alpha}$, where A is aggregate TFP and $N_C = \sum_{i \in \text{emp.}} z_{w,i}$ is corporate labor in efficiency units. No-arbitrage ($r_K = r$) pins the capital–labor ratio and the wage directly from the interest rate:

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

Entrepreneurs operate individual firms with decreasing-returns technology $y = \theta A z_e k^{\alpha_E}$, where θ is an entrepreneurial TFP parameter and k is capital. The unconstrained optimal scale is $k^* = (\alpha_E \theta A z_e / (r + \delta_k))^{1/(1-\alpha_E)}$; actual capital is bounded by the collateral constraint.

6.1.4 Collateral Constraints

A single entrepreneur (or an EE couple) may borrow up to a fraction λ_k of physical capital:

$$k \leq \frac{a}{1 - \lambda_k}. \tag{9}$$

In EE couples, household assets split optimally across two firms; no wage collateral is available.

One key feature of the model is the *spousal collateral channel*. In EW or WE couples (one entrepreneur, one worker), the constraint expands to include a fraction λ_w of the worker spouse's income:

$$k \leq \frac{a + \lambda_w \cdot \tilde{w}_{\text{spouse}} \cdot z_{w,\text{spouse}}}{1 - \lambda_k}, \tag{10}$$

where $\tilde{w}_{\text{spouse}} = w$ if the spouse is employed and $\tilde{w}_{\text{spouse}} = b_u w$ if unemployed. Collateral depends on the worker spouse's permanent type z_w (not z_e), so matching on z_w shapes borrowing capacity. When the worker spouse loses their job, pledgeable income drops from wz_w to $b_u wz_w$, tightening the constraint—the *collateral destruction channel*.

6.1.5 Employment Risk

Employment is a state variable $ie \in \{0, 1, 2\}$: employed ($ie = 0$), unemployed on UI ($ie = 1$), and incumbent entrepreneur ($ie = 2$). Workers transition via a Markov chain with separation rate σ_s and job-finding rate φ_f, yielding steady-state unemployment $\bar{u} = \sigma_s/(\sigma_s + \varphi_f)$. An employed worker earns wz_w; an unemployed worker receives $b_u wz_w$.

Incumbent entrepreneurs face an independent exit shock: with probability φ_E each period the firm is destroyed and the agent enters unemployment. With probability $1 - \varphi_E$ the firm continues.

In recession ($A < 1$), unemployment is countercyclical. The target rate rises to $u(A) = \bar{u} + u_{\text{sens}} \cdot \max(0, 1 - A)$, achieved by a common scaling factor s applied to both σ_s and φ_f:

$$\sigma_s(A) = s \cdot \sigma_s, \quad \varphi_f(A) = \varphi_f/s, \quad s = \sqrt{\frac{u(A)\,\varphi_f}{\sigma_s\,(1 - u(A))}}. \tag{11}$$

6.1.6 Occupational Choice

Singles. A single agent in state (a, z_e, ie) with permanent type z_w chooses between entrepreneurship (E) and corporate work (W) each period. The entrepreneur picks

capital k and savings a':

$$c + a' = (1 + r)a + \theta A z_e k^{\alpha_E} - (r + \delta_k)k, \qquad k \leq \frac{a}{1 - \lambda_k}, \quad a' \geq 0. \tag{12}$$

The worker consumes out of wage income: $c + a' = (1 + r)a + y_{\text{work}}(ie, z_w)$, where y_{work} equals wz_w if employed, $b_u wz_w$ if unemployed, or $\xi_E b_u wz_w$ for a voluntarily exiting entrepreneur. The agent chooses whichever occupation yields the higher value.

Couples. A couple in state $(a, z_{e,1}, z_{e,2}, ie_1, ie_2)$ with permanent types $(z_{w,1}, z_{w,2})$ selects from four configurations: WW (both work), EW (spouse 1 entrepreneur, spouse 2 worker), WE (symmetric), or EE (both entrepreneurs). In EW, the budget constraint is

$$c + a' = (1 + r)a + \theta A z_{e,1} k_1^{\alpha_E} - (r + \delta_k)k_1 + y_{\text{work}}(ie_2, z_{w,2}), \tag{13}$$

with the collateral constraint (10) using $z_{w,2}$. In EE, assets split optimally across two firms, each with its own z_e draw, and each spouse faces an independent exit probability φ_E. The couple selects the value-maximising configuration.

Recursive formulation. Because z_w is permanent and binary, the singles problem is solved separately for each $z_w \in \{z_w^L, z_w^H\}$ over the state (a, z_e, ie). The couples problem is solved separately for each $(z_{w,1}, z_{w,2}) \in \{z_w^L, z_w^H\}^2$ over $(a, z_{e,1}, z_{e,2}, ie_1, ie_2)$. For each state, the value function is

$$V = \max_{c \in \{\text{E,W}\}} V_c \quad \text{(singles)}, \qquad V = \max_{c \in \{\text{WW,EW,WE,EE}\}} V_c \quad \text{(couples)}, \tag{14}$$

where each option-specific value uses the Epstein–Zin aggregator with the appropriate continuation-value mixture over next-period employment and productivity transitions.

6.1.7 Equilibrium

A stationary equilibrium consists of a bond rate r^*, wage w^*, value functions, policy functions, and a stationary distribution μ such that: (i) policies are optimal given prices; (ii) w^* satisfies the no-arbitrage condition; (iii) the capital market clears, $K_C(r^*) + K_E =$

$\int a\, d\mu$; and (iv) μ is stationary under the induced transitions. Aggregation weights singles by $\Pr(z_w)$ and couples by the matching probabilities (5)–(7).

6.2 Solution and Calibration

6.2.1 Solution Method

The equilibrium interest rate is found by bisection on the capital-market excess demand. Each evaluation requires solving six VFI problems (two singles, four couples) and forward-iterating the associated stationary distributions. Howard policy-evaluation steps accelerate convergence in the presence of the discrete occupational choice. A policy-stable fallback declares convergence when the discrete policy has been bit-identical for eight iterations. Full algorithmic details are in Appendix B.

6.2.2 External Parameters

Table 9 reports all parameter values. The preference parameters—discount factor $\beta = 0.96$, risk aversion $\gamma = 3$, and IES $\psi = 0.50$—are within the standard range used in heterogeneous-agent macro models (Krusell and Smith, 1998; Cagetti and De Nardi, 2006; Quadrini, 2000). The corporate capital share $\alpha = 0.33$ and annual depreciation rate $\delta_k = 0.10$ are conventional. The separation rate $\sigma_s = 0.04$ and job-finding rate $\varphi_f = 0.95$ are chosen to match the Canadian unemployment rate of approximately 4% in the pre-oil-shock period, with an average unemployment spell of about one year. The UI replacement rate $b_u = 0.55$ reflects the Canadian Employment Insurance system; since self-employed business owners do not qualify for EI in Canada, exiting entrepreneurs receive $\xi_E = 0$.

Worker productivity types $z_w^L = 0.50$ and $z_w^H = 1.50$ represent workers below and above the median, with $p = 0.50$ splitting the population evenly. The ratio $z_w^H / z_w^L = 3$ approximately matches the ratio of average wages above versus below the median in Canadian earnings data. The fraction of individuals in couples is $m = 0.60$ from Canadian household surveys. Assortative matching is calibrated to $\rho_m = 0.30$ using educational homogamy data from the 2001–2006 Canadian Census: with $p = 0.50$, the conditional probability $P(H|H) = p + \rho_m(1-p) = 0.65$ matches the approximately 65%

same-education matching rate. The AR(1) process for entrepreneurial productivity z_e is parameterised with persistence $\rho_z = 0.90$ and volatility $\sigma_z = 0.20$, following Cagetti and De Nardi (2006) and Quadrini (2000), and discretised into $N_Z = 5$ grid points via Tauchen's method.

6.2.3 Internally Calibrated Parameters

Four parameters are determined to match Canadian data targets. The entrepreneurial capital curvature α_E, TFP coefficient θ, and capital collateral fraction λ_k are jointly calibrated via coarse grid search followed by Nelder–Mead refinement of a moment-matching loss function. The exit rate φ_E is pinned analytically.

The *entrepreneurial population share* target of approximately 15% is based on Canadian data: roughly 3.5 million private business owners and self-employed individuals out of a labor force of 22.2 million. The *average leverage* target of approximately 0.27 is the average debt-to-capital ratio among entrepreneurs observed in our matched treatment sample (Table 2 of the paper). The *aggregate capital-to-output ratio* target of approximately 3.0 corresponds to the Canadian aggregate K/Y from the Penn World Table. The *business exit rate* target of approximately 12% reflects Canadian business turnover data, and pins $\varphi_E = 0.12$ directly.

The calibration converges to $\alpha_E = 0.5585$, $\theta = 0.3701$, and $\lambda_k = 0.4760$.

6.2.4 Spousal Collateral

The wage collateral fraction $\lambda_w = 0.25$ is set externally rather than calibrated, because it governs the strength of the spousal collateral channel—the mechanism of interest. This parameter determines how much of the worker spouse's income the entrepreneur can pledge as additional collateral, and therefore controls the magnitude of both the collateral level advantage and the collateral destruction channel. Because the quantitative results are sensitive to this choice, we conduct a systematic sensitivity analysis in Section 6.3.4, varying λ_w from 0.10 to 0.50.

6.3 Quantitative Results

6.3.1 Untargeted Moments

Table 10 compares three moments not used in calibration against their empirical counterparts. The model closely matches the net-income-to-assets ratio $(0.104$ vs. $0.145)$ and the wealth Gini $(0.684$ vs. $0.690)$. The entrepreneur-to-worker earnings ratio is overstated $(2.499$ vs. $2.634)$, reflecting the binary z_w structure which generates wider cross-occupation earnings dispersion than observed in the data.

6.3.2 Difference-in-Differences: Disciplining the Spousal Channel

The model's central prediction is that married entrepreneurs respond differently to recessions than singles because spousal employment status affects household insurance and borrowing capacity. We test this prediction using a difference-in-differences framework applied to *incumbent* entrepreneurs—agents operating firms in the baseline steady state. For each economy, we hold the baseline distribution fixed and apply recession-era policy functions to the baseline population, isolating policy responses from distributional shifts. The DiD is

$$\Delta\Delta X = \left(\bar{X}^{\text{recess}}_{\text{couples}} - \bar{X}^{\text{base}}_{\text{couples}}\right) - \left(\bar{X}^{\text{recess}}_{\text{singles}} - \bar{X}^{\text{base}}_{\text{singles}}\right), \tag{15}$$

for $X \in \{K, \text{leverage}\}$.

We first estimate the DiD in the baseline economy, which features countercyclical unemployment risk. In this economy, all married entrepreneurs face the same aggregate recession and the same increase in spousal unemployment risk. We then partition incumbent married entrepreneurs by the realized employment outcome of the worker spouse. The *spouse keeps job* estimate conditions on couples in which the worker spouse remains employed during the recession. The *spouse loses job* estimate conditions on couples in which the worker spouse becomes unemployed. Because both groups are drawn from the same economy, the estimates share the same parameters, prices, aggregate shock, and unemployment process. Figure 7 reports these estimates.

The *spouse keeps job* estimate captures the response of entrepreneurs whose spouses

remain employed but whose spousal insurance has become less reliable. This case includes two forces. First, the worker spouse still provides household insurance and pledgeable income, which can support consumption and relax borrowing constraints relative to single entrepreneurs. Second, the rise in unemployment risk weakens the reliability of that insurance and can tighten expected borrowing capacity, leading married entrepreneurs to de-risk even before any job loss occurs. The estimated effects are negative: $\Delta\Delta K = -0.27$ pp and $\Delta\Delta\text{lev} = -0.51$ pp. Thus, even when the spouse keeps their job, married entrepreneurs reduce capital and leverage more than otherwise comparable singles.

The *spouse loses job* estimate captures the same recession and unemployment-risk environment, but for couples in which the risk is realized. When the worker spouse becomes unemployed, pledgeable income falls from wz_w to $b_u wz_w$, directly tightening the borrowing constraint and reducing household resources. The DiD therefore becomes more negative: $\Delta\Delta K = -0.66$ pp and $\Delta\Delta\text{lev} = -0.91$ pp. The difference between the *spouse loses job* and *spouse keeps job* estimates captures the additional effect of the realized job loss, equal to -0.39 pp for capital and -0.40 pp for leverage.

We then estimate a separate counterfactual economy with *no cyclical unemployment risk*. In this economy, the recession still lowers productivity ($A = 0.95$), but we set $u_{\text{sens}} = 0$, so the probability of spousal unemployment does not rise in downturns. This counterfactual removes the weakened-insurance channel and isolates the conventional stabilizing role of spousal income.

The results reverse sign. In the no-cyclical-unemployment economy, married entrepreneurs increase capital and leverage relative to single entrepreneurs during the recession: $\Delta\Delta K = +0.66$ pp and $\Delta\Delta\text{lev} = +0.09$ pp. Without elevated spousal unemployment risk, spousal income continues to provide insurance and collateral support in both good and bad times. Married entrepreneurs therefore retrench less than singles, consistent with the conventional intuition that spousal income stabilizes entrepreneurial activity.

Comparing the baseline economy to the no-cyclical-unemployment economy shows why countercyclical unemployment risk is central. When spousal unemployment risk does not rise, marriage stabilizes entrepreneurial firms during recessions. When un-

employment risk is countercyclical, however, the same household structure becomes a source of de-risking: even married entrepreneurs whose spouses remain employed reduce capital and leverage relative to singles. Thus, the model shows that it is not marriage per se that amplifies downturns, but the failure of spousal insurance precisely when aggregate conditions deteriorate.

6.3.3 Does Spousal Insurance Amplify Negative Productivity Shocks?

To quantify aggregate implications, we examine the steady-state output decline after a permanent TFP shock in our calibrated economy relative to a counterfactual economy in which all entrepreneurs and workers are single. We estimate the impact of the TFP shock in both economies and examine whether the married economy amplifies or mitigates the output decline. To do so, we solve four economies: (i) baseline ($m = 0.60$, $A = 1$), (ii) recession ($m = 0.60$, $A = 0.95$), (iii) all singles ($m = 0$, $A = 1$), and (iv) singles recession ($m = 0$, $A = 0.95$). All four economies share the same calibrated parameters and are solved independently in general equilibrium.

The baseline couples economy produces 1.2% more output than the all-singles economy ($Y = 2.054$ vs. 2.029), sustains a higher interest rate ($r^* = 0.0287$ vs. 0.0264), and channels a larger share of output through entrepreneurship (33.7% vs. 32.2%). These differences reflect two complementary advantages of household structure. First, the *collateral-level advantage*: when $\lambda_w > 0$, spousal income pledging directly relaxes borrowing constraints for EW/WE couples, allowing entrepreneurs to operate closer to their unconstrained scale. Second, even absent collateral pledging, spousal income provides implicit *consumption insurance* that reduces the need for precautionary savings, freeing a larger share of household wealth for deployment as productive capital. In steady state, married entrepreneurs thus optimally operate larger, more leveraged firms than observationally similar singles.

In recession, the couples economy suffers a larger output decline (-11.21% vs. -10.71%), implying an amplification of 4.7%. This amplification operates through two channels. When $\lambda_w > 0$, rising unemployment risk directly threatens pledgeable collateral, tightening borrowing constraints. Even without collateral pledging, however, the increase in spousal unemployment risk introduces uncertainty into the household

budget constraint: the entrepreneur anticipates that spousal job loss would reduce household income from wz_w to $b_u wz_w$, eroding the consumption insurance that supports larger firm scale in steady state. This triggers precautionary retrenchment—the entrepreneur scales back capital and rebuilds savings buffers—regardless of whether spousal income enters the collateral constraint. The amplification is thus the aggregate counterpart of the negative DiD documented in Section 6.3.2, reflecting both the collateral and budget-constraint channels.

6.3.4 How Important Is Pledgeable Spousal Income?

We next assess how much of the model's amplification comes from the pledgeability of spousal income. The parameter λ_w governs the fraction of spousal labor income that entrepreneurs can pledge as collateral. If the spousal collateral channel is central to the amplification mechanism, changes in λ_w should materially affect the model's aggregate response to a negative productivity shock. For each value of λ_w, we re-solve the economy in general equilibrium.

Table 13 reports the results. Aggregate amplification changes little across values of λ_w, ranging from 4.6% to 4.9%. Even when λ_w is close to zero, so that spousal income has almost no collateral value, the model still generates large amplification. This result suggests that pledgeable spousal income is not the main driver of the aggregate effect. Instead, amplification is largely driven by the insurance role of spousal income: when spousal unemployment risk rises in downturns, entrepreneurs respond to the weakening of household insurance even if spousal income does little to relax borrowing constraints.

The DiD estimates also remain sizable when λ_w is low, again pointing to the importance of spousal income insurance rather than collateral alone. Pledgeable income becomes quantitatively important only at very high values of λ_w: when 50% of spousal income is pledgeable, the spouse-loses-job DiD becomes substantially more negative. Thus, spousal collateral can amplify the response of directly affected entrepreneurs when pledgeability is unreasonably high.

7 Conclusion

Entrepreneurial firms are typically owned and operated by individuals who are part of multi-person households. As a result, firm decisions are made in the presence of household income risk. We show that this household dimension matters for firm policy. Spousal labor income can insure household consumption and relax financial constraints, but this insurance is itself risky: it may weaken precisely when the entrepreneur's firm is most vulnerable.

We provide reduced-form evidence using the unexpected collapse in global oil prices in 2014 as a quasi-experimental increase in spousal unemployment risk. Using Canadian administrative data that link entrepreneurs to their spouses and to detailed firm balance sheets, we show that entrepreneurs adjust investment, leverage, and payout policies when spousal income insurance becomes less reliable. The patterns are consistent with precautionary responses to weaker household insurance: the effects are stronger among entrepreneurs with characteristics associated with higher risk aversion and among firms with riskier cash flows.

We also show that this household channel has aggregate implications. In a calibrated heterogeneous-agent general-equilibrium model, spousal income supports entrepreneurial activity in normal times, but countercyclical unemployment risk weakens this support during downturns. As a result, marriage can stabilize entrepreneurial firms in good times while amplifying output losses in bad times.

Overall, the paper shows that entrepreneurs do not make firm decisions as standalone agents. They internalize risks that originate elsewhere in the household and adjust firm policies when household insurance becomes less reliable. This household-to-firm channel provides a mechanism through which labor-market risk can shape firm behavior and amplify aggregate downturns.

References

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

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.

Andersen, Steffen, and Kasper Meisner Nielsen, 2012, Ability or finances as constraints on entrepreneurship? evidence from survival rates in a natural experiment, *The Review of Financial Studies* 25, 3684–3710.

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.

Bianchi, Milo, and Matteo Bobba, 2013, Liquidity, risk, and occupational choices, *Review of Economic Studies* 80, 491–511.

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.

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

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, Sup-

ply 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.

Chen, Hui, Jianjun Miao, and Neng Wang, 2010, Entrepreneurial finance and nondiversifiable risk, *The Review of Financial Studies* 23, 4348–4388.

Chen, Zhiyao, and Ilya A Strebulaev, 2019, Macroeconomic risk and idiosyncratic risk-taking, *The Review of Financial Studies* 32, 1148–1187.

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.

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

De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo, 2025, Wage risk and government and spousal insurance, *Review of Economic Studies* 92, 954–980.

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.

Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri, 2018, Portfolio choices, firm shocks, and uninsurable wage risk, *The Review of Economic Studies* 85, 437–474.

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.

Gottlieb, Joshua D, Richard R Townsend, and Ting Xu, 2022, Does career risk deter potential entrepreneurs?, *The Review of Financial Studies* 35, 3973–4015.

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.

Holtz-Eakin, Douglas, David Joulfaian, and Harvey S Rosen, 1993, Entrepreneurial decisions and liquidity constraints.

Hombert, Johan, Antoinette Schoar, David Sraer, and David Thesmar, 2020, Can unemployment insurance spur entrepreneurial activity? evidence from france, *The Journal of Finance* 75, 1247–1285.

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).

Jahn, George, 2014, Opec decides not to cut oil production, sending crude prices to 4-year low, *The Washington Post* Associated Press.

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* .

Krauss, Clifford, 2014, Despite price slump, no end in sight for oil boom, *The State* New York Times News Service. Updated October 20, 2014.

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.

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

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.

Panousi, Vasia, and Dimitris Papanikolaou, 2012, Investment, idiosyncratic risk, and ownership, *The Journal of Finance* 67, 1113–1148.

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.

Silva, Dejanir, and Robert Townsend, 2025, Risk-taking over the life cycle: Aggregate and distributive implications of entrepreneurial risk, *manuscript, University of Illinois at Urbana-Champaign* 268–316.

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 1: Share of married or cohabiting entrepreneurs

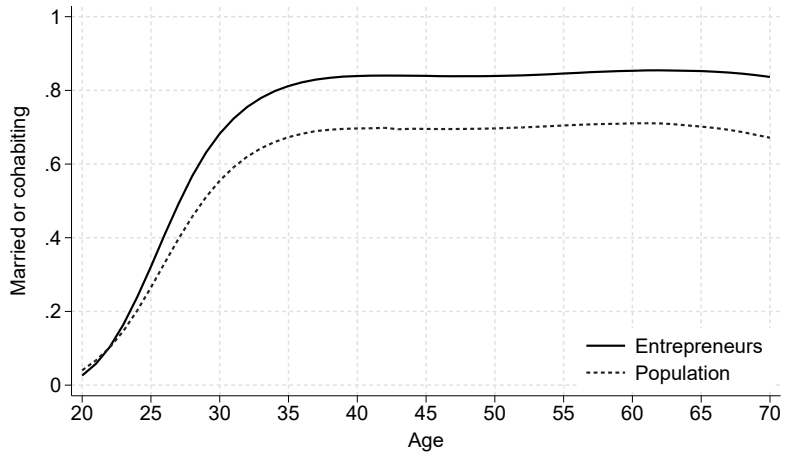

Notes: This figure plots the share of individuals who are married or cohabiting at each age for the universe of entrepreneurs and for the overall population in Canada between 2001 and 2017.

Figure 2: Oil prices and macroeconomic indicators in Canada

(**a**) Oil price and GDP per capita

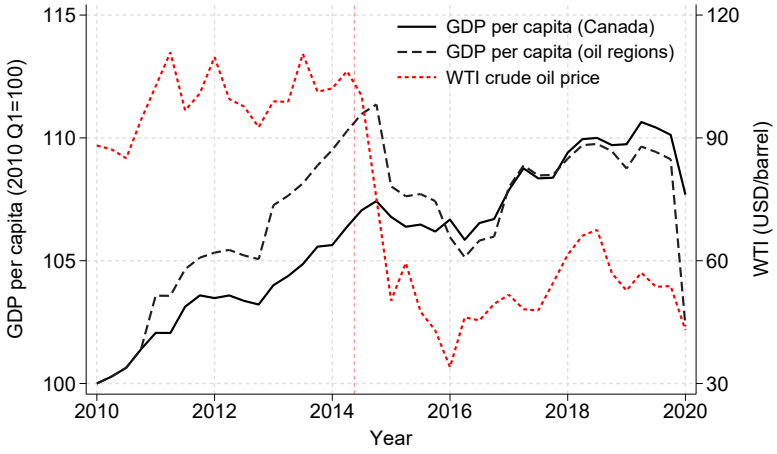

(**b**) Oil price and unemployment

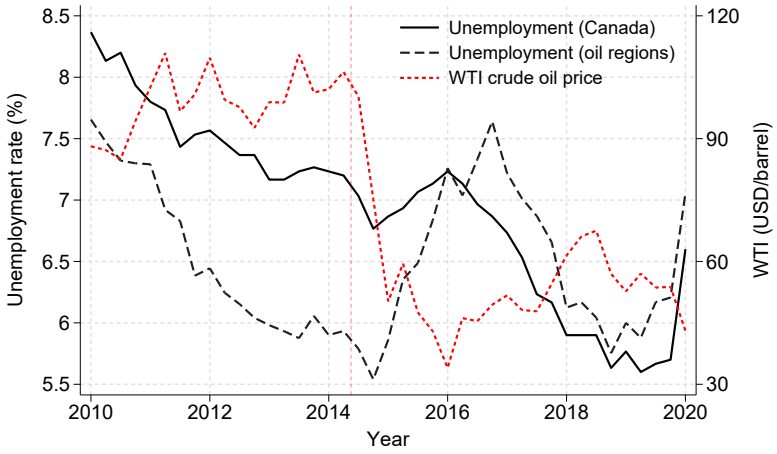

Notes: This figure presents the relationship between oil prices and aggregate macroeconomic conditions in Canada and in the main oil-exposed provinces (Alberta, Saskatchewan, British Columbia, and Newfoundland and Labrador). Panel A plots quarterly WTI crude oil prices (in 2017 USD) against growth in 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 3: Unemployment rates across Canadian Economic Regions

(**a**) Unemployment rates in 2013

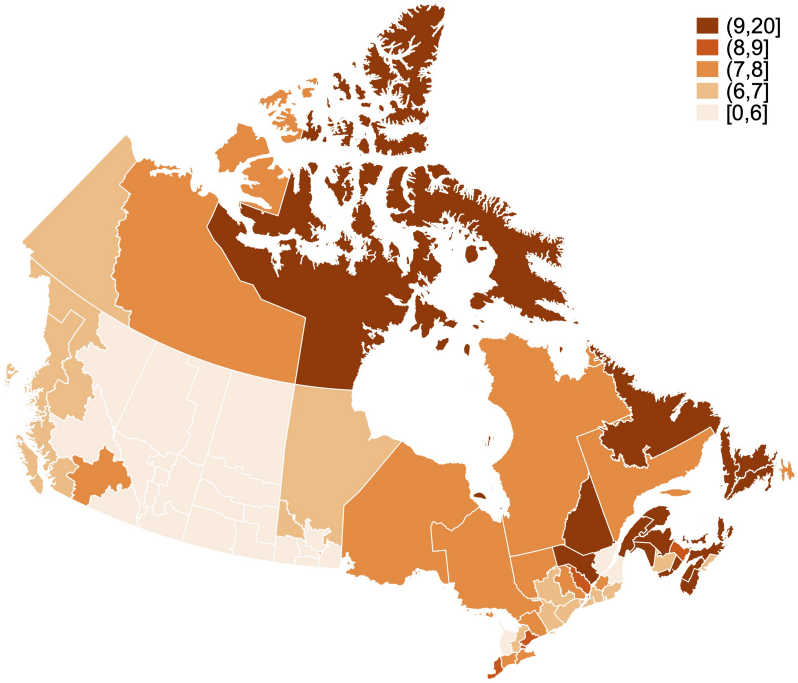

(**b**) Unemployment rates in 2016

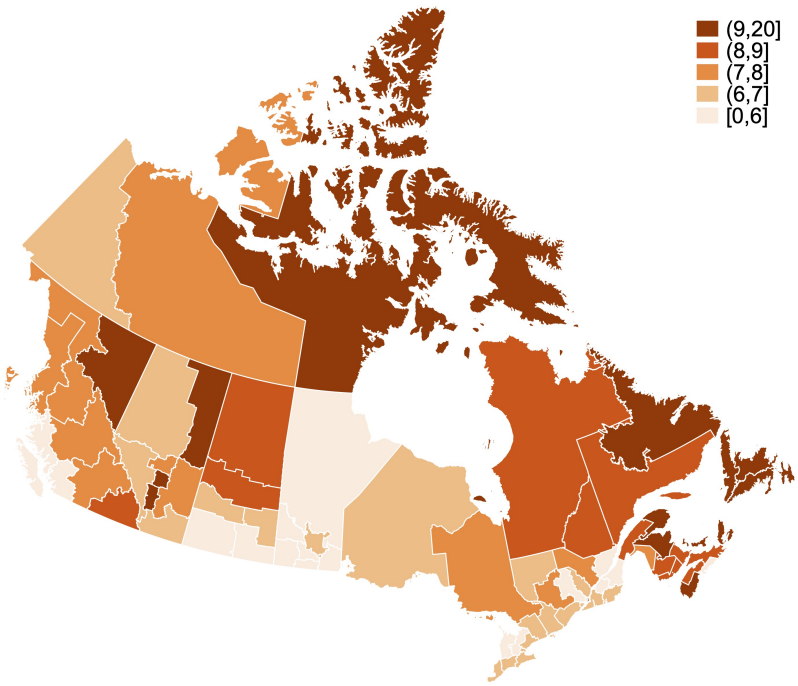

Notes: This figure shows unemployment rates across Canadian Economic Regions in 2013 (Panel A) and 2016 (Panel B). Economic Regions are defined by Statistics Canada as standard geographic units for analysis of regional economic activity. An Economic Region contains multiple commuting zones. Unemployment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

Figure 4: Covariate balance

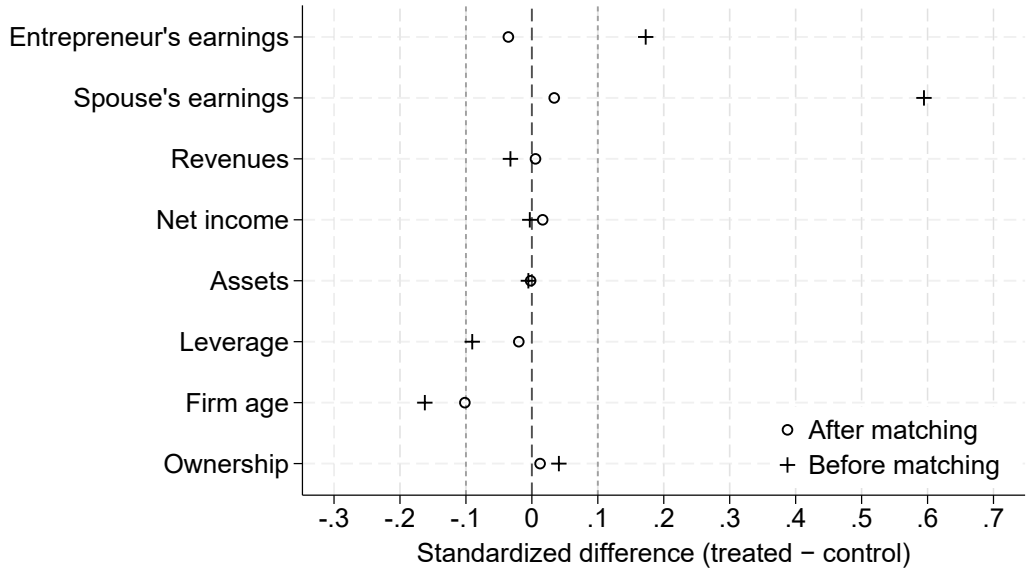

Notes: This figure reports the standardized difference in means between treated and control firms for the continuous variables listed in Table 1. Treated firms are those whose owner's spouse was employed in the oil sector in 2013 and 2014. Potential controls are firms owned by entrepreneurs with spouses 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 5: Effects of the oil shock on spouse's job loss and unemployment risk

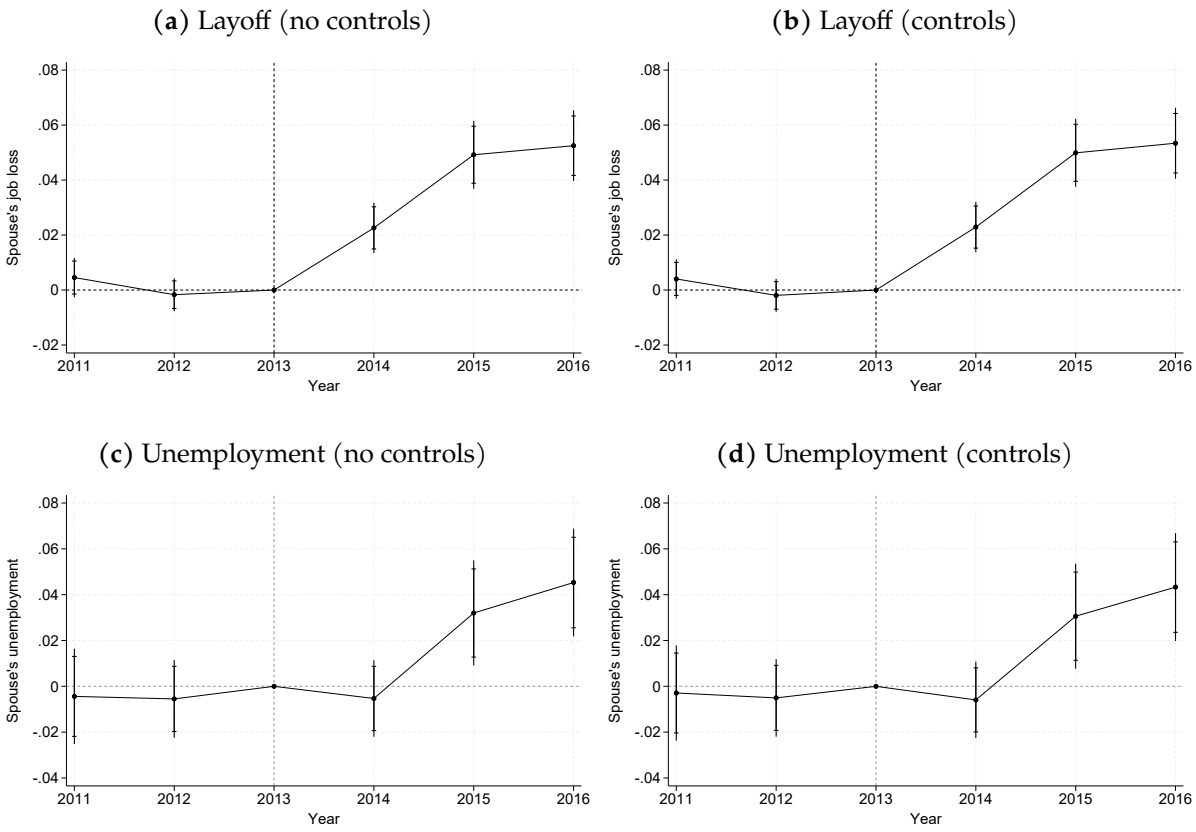

Notes: This figure shows event–study estimates from equation 1 on an indicator indicator for whether an individual experiences a layoff (panels (a)–(b)) or unemployment (panels (c)–(d)). Layoff is an indicator equal to 1 when a worker's employment spell ends and the employer records the reason for the 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 6: Spouse's unemployment risk and firm outcomes

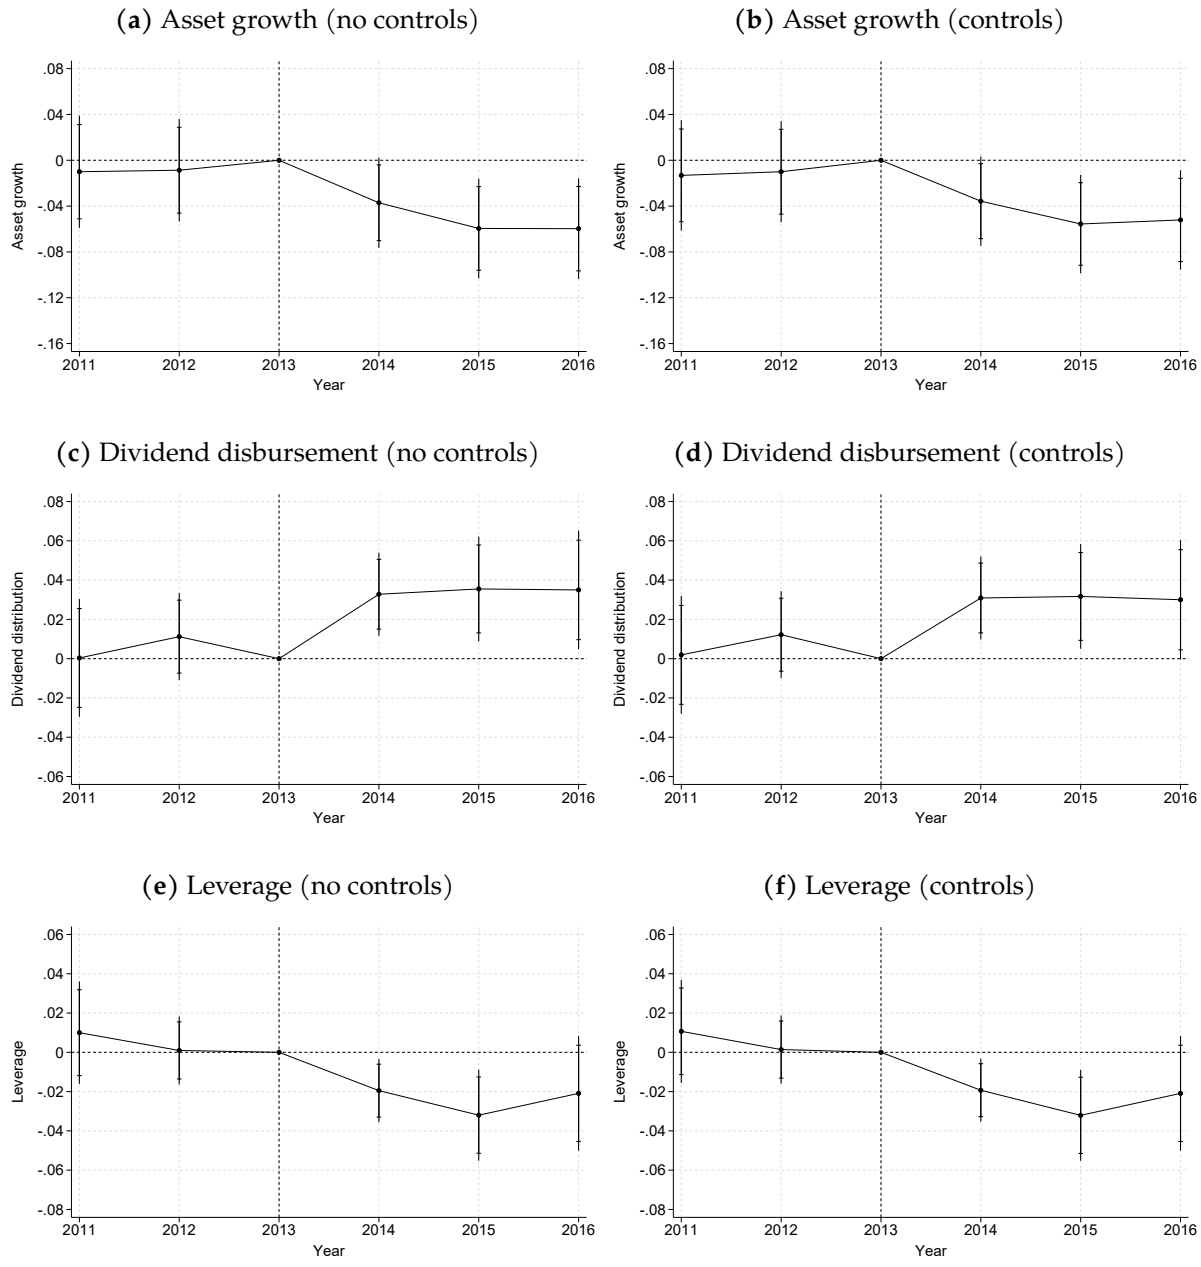

Notes: This figure reports event–study estimates obtained by fitting equation 2 on asset growth, dividend disbursement, and financial leverage. The treated group consists of entrepreneurs with spouses employed in the oil sector in both 2013 and 2014, while the matched control group includes entrepreneurs 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 7: Model Difference-in-Differences

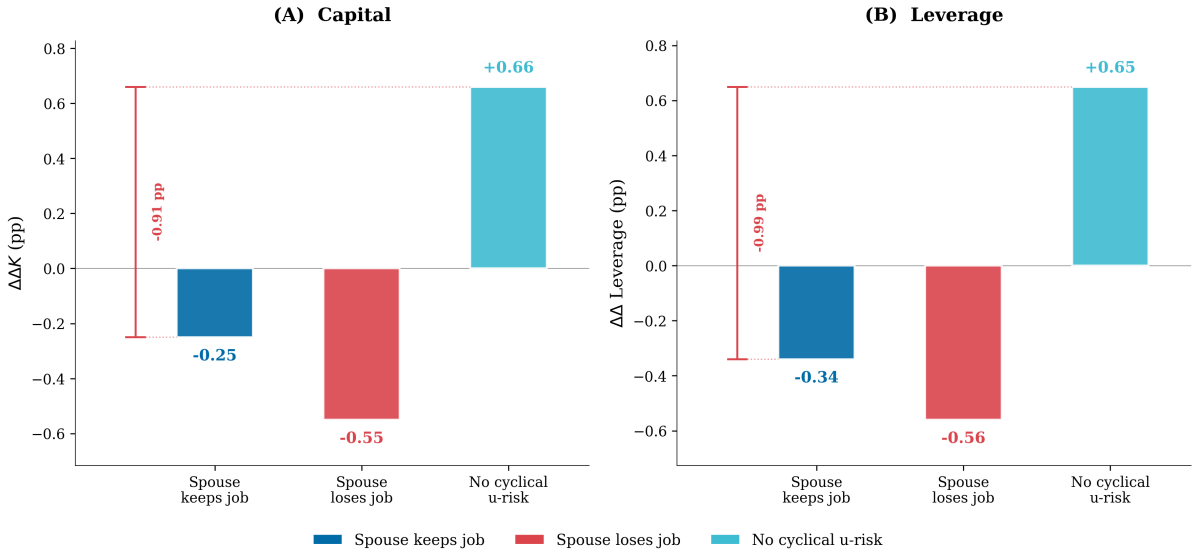

Notes: Difference-in-differences by scenario: capital (Panel A) and leverage (Panel B). The bracket between "Spouse keeps job" and "No cyclical u-risk" isolates the precautionary savings motive due to increased unemployment risk.

Figure 8: Sensitivity to Spousal Income Pledgeability (λ_w)

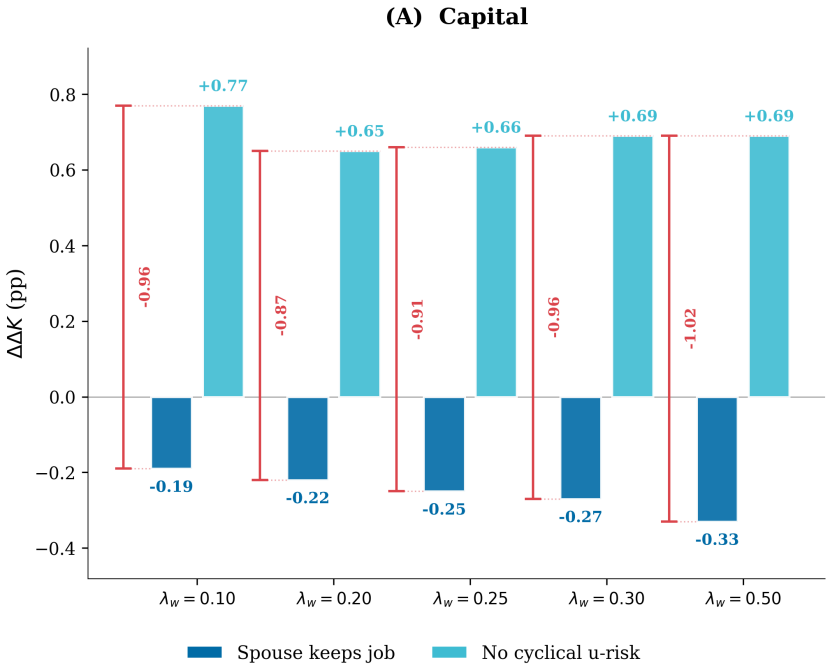

Notes: DiD for capital across λ_w values. Each group shows the "Spouse keeps job" and "No cyclical u-risk" scenarios; the bracket isolates the precautionary savings motive. As λ_w rises, the precautionary effect grows while the collateral level advantage remains stable.

Tables

Table 1: Summary statistics

	Matched treated (1)	Matched control (2)	All treated (3)	Potential control (4)
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)	1,865.67	1,953.33	2,039.87	1,460.48
	(2,109.55)	(2,580.73)	(2,676.33)	(1,922.74)
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 2: 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 3: 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 3 (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 4: Heterogeneous effects on investment by entrepreneur's share of household income

	Low HH income share		Middle HH income share		High HH income share	
	(1)	(2)	(3)	(4)	(5)	(6)
Post × $\text{Oil}_{i(j)}$	-0.0396	-0.0279	-0.0765***	-0.0631***	-0.00475	-0.0159
	(0.0296)	(0.0291)	(0.0227)	(0.0225)	(0.0453)	(0.0451)
Controls	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × commuting zone × year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6,295	6,295	11,115	11,115	2,730	2,730
R^2	0.412	0.430	0.427	0.444	0.455	0.470

Notes: This table reports heterogeneity in treatment effects on asset growth by the entrepreneur's share of household income. 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 $\Delta a_t = \log(a_t) - \log(a_{t-1})$. Low income share refers to entrepreneurs who represent at most one-third of household income in 2013, middle income share to entrepreneurs who represent between one-third and two-thirds, and high income share to entrepreneurs who represent more than two-thirds. 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 5: Heterogeneous effects on investment by household precautionary motives

Panel A. By entrepreneur age and gender

	(1)	(2)	(3)	(4)
Post $\times$ Oil$_{i(j)}$	-0.033**	-0.030**	-0.041***	-0.036***
	(0.016)	(0.015)	(0.014)	(0.014)
Post $\times$ Oil$_{i(j)}$ $\times$ $\mathbb{I}$\{Age $>$ p75\}	-0.066*	-0.067*		
	(0.037)	(0.036)		
Post $\times$ Oil$_{i(j)}$ $\times$ $\mathbb{I}$\{Female, Age $>$ p75\}			-0.091*	-0.094*
			(0.054)	(0.052)
Controls	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry $\times$ commuting zone $\times$ year FE	Yes	Yes	Yes	Yes
Observations	22,630	22,630	22,630	22,630
R^2	0.421	0.435	0.421	0.435

Table 5 (continued)

Panel B. By parental status

	No young children		At least 1 child		At least 2 children	
	(1)	(2)	(3)	(4)	(5)	(6)
Post × $\text{Oil}_{i(j)}$	-0.048***	-0.040**	-0.072*	-0.068*	-0.116***	-0.106**
	(0.017)	(0.017)	(0.039)	(0.038)	(0.044)	(0.042)
Controls	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × commuting zone × year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17,415	17,415	4,445	4,445	3,480	3,480
R^2	0.415	0.431	0.424	0.441	0.438	0.457

Notes: This table reports heterogeneous treatment effects on asset growth by household precautionary motives. 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 whose spouses were employed in the oil sector in both 2013 and 2014, while the matched control group includes entrepreneurs whose spouses worked in other sectors. The dependent variable is asset growth, defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. Panel A examines heterogeneity by entrepreneur age and gender. Panel A restricts the sample to couples whose age difference does not exceed 20 years. $\mathbb{I}\{\text{Age} > \text{p75}\}$ identifies older couples, defined as those in which both the entrepreneur and the spouse are above the 75th percentile of their respective age distributions. $\mathbb{I}\{\text{Female, Age} > \text{p75}\}$ identifies older couples satisfying the same age criteria in which the entrepreneur is female. Panel B examines heterogeneity by the presence of children in the household. Columns (1)–(2) restrict the sample to entrepreneurs with no young children, where young children are defined as children younger than 8 in 2014. Columns (3)–(4) restrict the sample to entrepreneurs with at least one child, and columns (5)–(6) to entrepreneurs with at least two children. 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: 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 7: Placebo test using spouses who leave the oil industry before the shock

	18-month window		12-month window		6-month window	
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Oil leaver$_{i(j)}$	0.0251	0.0429	-0.00572	0.0178	0.0140	0.0280
	(0.0384)	(0.0369)	(0.0419)	(0.0409)	(0.0490)	(0.0530)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year FE	Yes	No	Yes	No	Yes	No
Commuting zone × year FE	Yes	No	Yes	No	Yes	No
Industry × commuting zone × year FE	No	Yes	No	Yes	No	Yes
Observations	4,890	4,890	3,880	3,880	2,625	2,545
R^2	0.417	0.436	0.396	0.420	0.419	0.457

Notes: This table reports placebo estimates for entrepreneurial firms based on equation 2. The outcome is asset growth, defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. Post is equal to 1 in 2014 or after. Oil leaver$_{i(j)}$ is an indicator equal to 1 for entrepreneurs whose spouses leave employment in the oil sector before the 2014 oil price collapse and move to a different sector. The treated group consists of entrepreneurs whose spouses were employed in the oil sector before the shock but leave the sector prior to the collapse, while the matched control group includes entrepreneurs whose spouses worked in other sectors and also leave their sector over the same period. Columns (1)–(2) use the broadest definition, including spouses who leave the oil industry between January 2013 and June 2014. Columns (3)–(4) restrict attention to spouses who leave within the 12 months preceding the shock. Columns (5)–(6) use the most restrictive definition, including only spouses who leave in the first half of 2014. 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 8: 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.

Table 9: Parameter Values

Parameter	Symbol	Value	Source / Target
External: Preferences			
Discount factor	β	0.96	Standard annual
Risk aversion	γ	3	Macro literature
IES	ψ	0.50	Macro-finance literature
External: Technology			
Corporate capital share	α	0.33	Standard
Depreciation rate	δ_k	0.10	Annual
External: Labor market			
Separation rate	σ_s	0.04	Canadian $u \approx 4\%$ (pre-oil-shock)
Job-finding rate	φ_f	0.95	Avg. duration ≈ 1 year
UI replacement rate	b_u	0.55	Canadian EI system
UI for exiting entrepreneurs	ξ_E	0.00	Self-employed ineligible for EI
External: Population and matching			
Fraction in couples	m	0.60	Canadian household data
Worker types	z_w^L, z_w^H	$0.50, 1.50$	Above/below median; ratio 3 (hours-adj.)
Fraction high type	p	0.50	Median split
Matching correlation	ρ_m	0.30	Canadian educ. homogamy (65%)
External: Entrepreneurial productivity			
Persistence	ρ_z	0.90	Cagetti & De Nardi (2006)
Volatility	σ_z	0.20	Cagetti & De Nardi (2006)
External: Spousal collateral			
Wage collateral	λ_w	0.25	Mechanism of interest
Internally calibrated			
Capital curvature	α_E	0.5880	$\rightarrow$ Entrep. share $= 16.1\%$
TFP coefficient	θ	0.3470	$\rightarrow K/Y = 3.2$ (PWT, non-residential)
Capital collateral	λ_k	0.6540	$\rightarrow$ Avg. leverage $= 0.342$
Exit rate	φ_E	0.08	$\rightarrow$ Exit rate $= 8\%$

Notes: External parameters are set from the macro literature or Canadian data. Internally calibrated parameters are jointly determined to match Canadian entrepreneurship moments.

Table 10: Untargeted Moments

Moment	Data	Model	Diff.
Net Income / Assets (couples)	0.145	0.104	-0.041
Wealth Gini	0.690	0.684	-0.006
E earnings / W earnings (couples)	2.634	2.499	-0.135

Notes: Moments not targeted in calibration, computed at the baseline economy. Net income and earnings are averages across EW/WE couples in the stationary distribution. The wealth Gini is household-level. Data: matched treated sample (Table 3 of the paper) and the Survey of Financial Security.

Table 11: DiD: Couples − Singles Incumbent Response

Scenario	$\Delta\Delta K$ (pp)	$\Delta\Delta$ lev (pp)
Spouse keeps job	-0.25	-0.34
Spouse loses job	-0.55	-0.56
No cyclical u-risk	$+0.66$	$+0.65$

Notes: Difference-in-differences for incumbent entrepreneurs. Couples are the treatment group; singles are the control. Recession: $A = 0.95$, unemployment rises from 4% to 9%. No incumbent voluntarily exits in the recession with cyclical unemployment, so these estimates are free of selective exit.

Table 12: Counterfactual Economies

Economy	A	r^*	w^*	Y	ent%	lev	K/Y	Y_E/Y
Baseline	1.00	0.0296	1.062	2.219	15.9%	0.337	3.22	39.3%
Recession	0.95	0.0279	0.990	1.967	15.9%	0.346	3.27	39.7%
All singles	1.00	0.0273	1.071	2.201	15.8%	0.370	3.23	38.1%
Singles recession	0.95	0.0239	1.006	1.968	15.9%	0.381	3.32	38.7%

Notes: Four economies solved independently in general equilibrium. Amplification: couples economy output decline (-11.36%) exceeds singles (-10.58%) by 7.4%.

Table 13: Sensitivity to λ_w: Amplification and DiD ($\Delta\Delta K$, pp)

	$\lambda_w = 0.10$	$\lambda_w = 0.20$	$\lambda_w = 0.25$	$\lambda_w = 0.30$	$\lambda_w = 0.50$
Spouse keeps job	-0.19	-0.22	-0.25	-0.27	-0.33
Spouse loses job	-0.31	-0.47	-0.55	-0.63	-0.93
No cyclical u-risk	$+0.77$	$+0.65$	$+0.66$	$+0.69$	$+0.69$

Notes: All economies re-solved in general equilibrium for each λ_w value; $\lambda_w = 0.25$ is the calibrated baseline.

A Additional Figures and Tables

Figure A.1: Firm policies with no realized negative income shocks

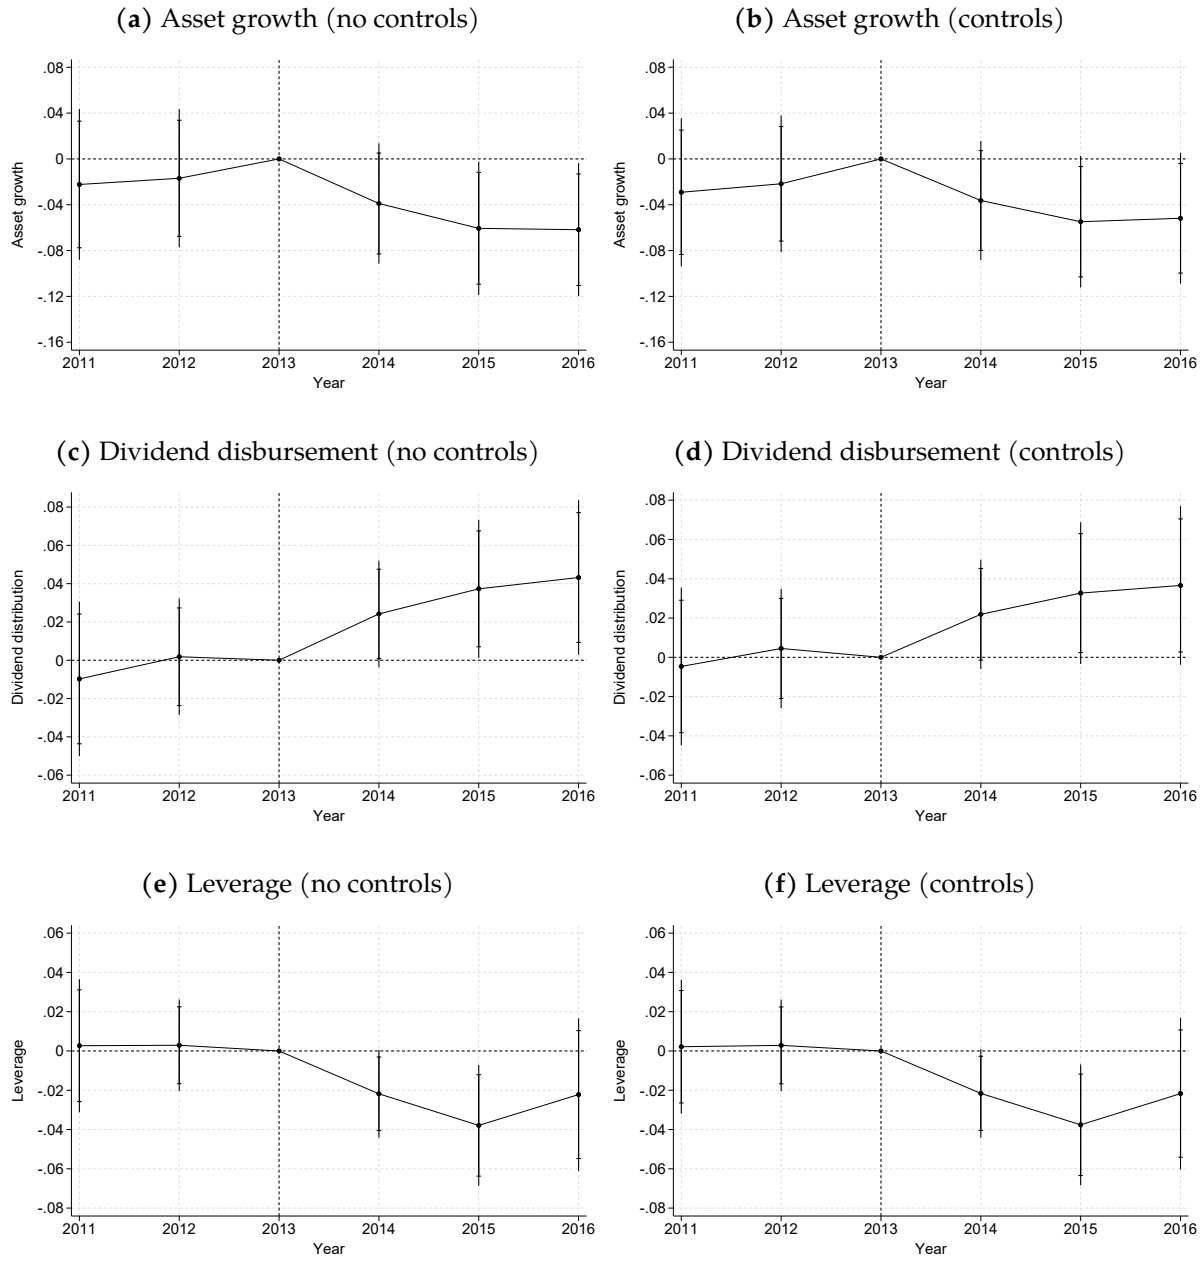

Notes: This figure reports event–study estimates obtained by fitting equation 2 on on asset growth, dividend disbursement, and financial leverage. We restrict the sample to matched treated-control pairs in which neither spouse experiences a large negative realized income shock during the outcome window, defined as earnings growth in the bottom quartile of the distribution. 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.

Figure A.2: Spouse's unemployment risk and firm profitability

(**a**) Profit margin (no controls)

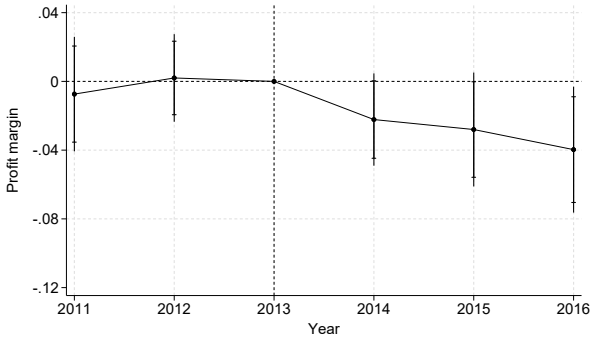

(**b**) Profit margin (controls)

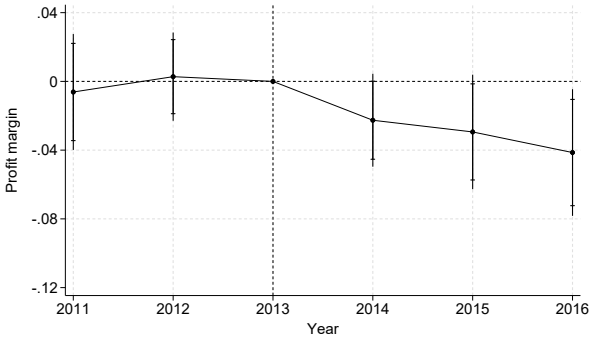

*Notes:*This figure reports event–study estimates obtained by fitting equation 2 on profit margin. The treated group consists of entrepreneurs with spouses employed in the oil sector in both 2013 and 2014, while the matched control group includes entrepreneurs with spouses employed in other sectors. 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.

Table A.1: Spouse's oil-sector attachment and firm investment

	Low oil-sector attachment				High oil-sector attachment			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post × $\text{Oil}_{i(j)}$	-0.032	-0.014	-0.039	-0.019	-0.058***	-0.056***	-0.058***	-0.058***
	(0.023)	(0.022)	(0.024)	(0.024)	(0.018)	(0.017)	(0.018)	(0.018)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year FE	Yes	Yes	No	No	Yes	Yes	No	No
Commuting zone × year FE	Yes	Yes	No	No	Yes	Yes	No	No
Industry × commuting zone × year FE	No	No	Yes	Yes	No	No	Yes	Yes
Observations	10,390	10,390	9,345	9,345	14,970	14,970	14,120	14,120
R^2	0.379	0.395	0.468	0.482	0.347	0.367	0.415	0.432

Notes: This table reports heterogeneity in the treatment effect on firm investment by the spouse's pre-shock attachment to the oil sector. Firm investment is measured by asset growth, defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. For treated entrepreneurs, oil-sector attachment is measured as the share of the spouse's observed career, from first entry into the labor market through the pre-shock period, spent working in the oil sector. Columns (1)–(4) report estimates for spouses who spent less than half of their observed career in the oil sector, while columns (5)–(8) report estimates for spouses who spent at least half of their observed career in the oil sector. Greater oil-sector attachment proxies for a higher cost of job displacement following the oil shock. Post is equal to 1 in 2014 or after. $\text{Oil}_{i(j)}$ is an indicator equal to 1 for entrepreneurs whose spouses were employed in the oil sector in both 2013 and 2014, while 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: Spouse's industry exposure to the oil shock and firm investment

Panel A. Control spouses in high oil-beta industries

	(1)	(2)	(3)	(4)
Post × $\text{Oil}_{i(j)}$	-0.010	-0.011	-0.013	-0.011
	(0.023)	(0.022)	(0.024)	(0.023)
Controls	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry × year FE	Yes	Yes	No	No
Commuting zone × year FE	Yes	Yes	No	No
Industry × commuting zone × year FE	No	No	Yes	Yes
Observations	9,100	9,100	8,375	8,375
R^2	0.380	0.406	0.455	0.476

Panel B. Control spouses in low oil-beta industries

	(1)	(2)	(3)	(4)
Post × $\text{Oil}_{i(j)}$	-0.054***	-0.052**	-0.058***	-0.053**
	(0.021)	(0.020)	(0.022)	(0.022)
Controls	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry × year FE	Yes	Yes	No	No
Commuting zone × year FE	Yes	Yes	No	No
Industry × commuting zone × year FE	No	No	Yes	Yes
Observations	11,825	11,825	10,995	10,995
R^2	0.359	0.375	0.433	0.447

Notes: This table examines how the estimated effect on firm investment varies with the oil exposure of the industries in which control spouses are employed. Firm investment is measured by asset growth, defined as $\Delta a_t = \log(a_t) - \log(a_{t-1})$. Oil exposure is measured using the beta of the spouse's industry with respect to the oil sector. Panel A restricts the control group to spouses employed in industries with high oil beta, while Panel B restricts the control group to spouses employed in industries with low oil beta. The treated group consists of entrepreneurs whose spouses were employed in the oil sector in both 2013 and 2014. Post is equal to 1 in 2014 or after. 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: Spouse's unemployment risk and entrepreneur's payout

	Total compensation				Salary				Dividends			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Post × $\text{Oil}_{i(j)}$	0.051**	0.056***	0.047**	0.054**	0.009*	0.009*	0.007	0.007	0.039*	0.041**	0.040**	0.043**
	(0.021)	(0.021)	(0.021)	(0.021)	(0.005)	(0.005)	(0.006)	(0.006)	(0.020)	(0.020)	(0.020)	(0.020)
Controls	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ind. × year FE	Yes	Yes	No	No	Yes	Yes	No	No	Yes	Yes	No	No
CZ × year FE	Yes	Yes	No	No	Yes	Yes	No	No	Yes	Yes	No	No
Ind. × CZ × year FE	No	No	Yes	Yes	No	No	Yes	Yes	No	No	Yes	Yes
Observations	23,070	23,070	21,405	21,405	23,070	23,070	21,405	21,405	23,070	23,070	21,405	21,405
R^2	0.601	0.602	0.626	0.627	0.731	0.732	0.757	0.758	0.608	0.608	0.631	0.632

Notes: This table reports average treatment effects from equation 2 on distributions from the firm to the entrepreneur. 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. Columns (1)–(4) examine total compensation paid to the entrepreneur, scaled by firm sales. Columns (5)–(8) and (9)–(12) decompose total compensation into its salary and dividend components, respectively, with each component also scaled by firm sales. 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.4: Spouse's unemployment risk and operating performance

	Sales growth				Profit margin			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post × $\text{Oil}_{i(j)}$	-0.021	-0.014	-0.027	-0.018	-0.024**	-0.027**	-0.025**	-0.028**
	(0.019)	(0.018)	(0.020)	(0.019)	(0.012)	(0.012)	(0.012)	(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	No	No	Yes	Yes	No	No
Commuting zone × year FE	Yes	Yes	No	No	Yes	Yes	No	No
Industry × commuting zone × year FE	No	No	Yes	Yes	No	No	Yes	Yes
Observations	21,190	21,190	19,450	19,450	26,205	26,205	24,715	24,715
R^2	0.350	0.371	0.428	0.445	0.597	0.598	0.646	0.646

Notes: This table reports average treatment effects from equation 2 on operating performance. Columns (1)–(4) examine sales growth. Columns (5)–(8) examine profit margin. 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 whose spouses were employed in the oil sector in both 2013 and 2014, while 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.5: Heterogeneous effects on investment by firm characteristics

Panel A. By firm size

	Below median				Above median			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post × $\text{Oil}_{i(j)}$	-0.044*	-0.038	-0.070**	-0.062**	-0.054***	-0.047***	-0.074***	-0.066***
	(0.025)	(0.025)	(0.028)	(0.027)	(0.015)	(0.015)	(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	No	No	Yes	Yes	No	No
Commuting zone × year FE	Yes	Yes	No	No	Yes	Yes	No	No
Industry × commuting zone × year FE	No	No	Yes	Yes	No	No	Yes	Yes
Observations	11,780	11,780	10,055	10,055	13,265	13,260	11,320	11,315
R^2	0.328	0.347	0.392	0.409	0.405	0.424	0.480	0.496

Table A.5 (continued)

Panel B. By firm profitability

	Below median				Above median			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post × $\text{Oil}_{i(j)}$	-0.024	-0.012	-0.035	-0.024	-0.071***	-0.071***	-0.078***	-0.070***
	(0.019)	(0.019)	(0.022)	(0.022)	(0.021)	(0.020)	(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	No	No	Yes	Yes	No	No
Commuting zone × year FE	Yes	Yes	No	No	Yes	Yes	No	No
Industry × commuting zone × year FE	No	No	Yes	Yes	No	No	Yes	Yes
Observations	12,645	12,645	10,675	10,675	12,210	12,210	10,550	10,550
R^2	0.361	0.347	0.431	0.418	0.375	0.353	0.426	0.405

Notes: This table reports heterogeneity in the treatment effects on asset growth by pre-shock firm characteristics. Panel A splits firms by firm size in 2013. Columns (1)–(4) restrict the sample to firms below the median of the size distribution, while columns (5)–(8) restrict the sample to firms above the median. Panel B analogously splits firms by profitability in 2013, with columns (1)–(4) corresponding to firms below the median and columns (5)–(8) to firms above the median. Post is equal to 1 in 2014 or after. $\text{Oil}_{i(j)}$ is an indicator equal to 1 for entrepreneurs whose spouses were employed in the oil sector in both 2013 and 2014, while the matched control group includes entrepreneurs whose spouses worked in other sectors. Asset growth is 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.6: Unemployment risk and investment for continuously married households

	(1)	(2)	(3)	(4)	(5)	(6)
Post $\times$ $\text{Oil}_{i(j)}$	-0.041***	-0.033**	-0.041***	-0.033**	-0.049***	-0.040***
	(0.014)	(0.013)	(0.014)	(0.014)	(0.014)	(0.014)
Controls	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry $\times$ year FE	Yes	Yes	Yes	Yes	No	No
Commuting zone $\times$ year FE	No	No	Yes	Yes	No	No
Industry $\times$ commuting zone $\times$ year FE	No	No	No	No	Yes	Yes
Observations	25,355	25,355	25,235	25,235	23,625	23,625
R^2	0.315	0.335	0.334	0.352	0.410	0.427

Notes: This table reports average treatment effects from equation 2 on asset growth for the sample of couples who remain married throughout the post-shock period. Asset growth is 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. $\text{Oil}_{i(j)}$ is an indicator equal to 1 for entrepreneurs whose spouses were employed in the oil sector in both 2013 and 2014, while 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.7: Spouse's unemployment risk and firm investment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Layoff rate	-0.075***	-0.072***	-0.074***	-0.071***	-0.065***	-0.062***	-0.059***	-0.057**
	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
Spouse laid off		-0.023***		-0.021***		-0.019***		-0.018***
		(0.003)		(0.003)		(0.003)		(0.003)
Log spouse earnings		0.016***		0.015***		0.015***		0.015***
		(0.001)		(0.001)		(0.001)		(0.001)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	No	No	No	No	No	No
Industry $\times$ year FE	No	No	Yes	Yes	Yes	Yes	No	No
Province $\times$ year FE	No	No	No	No	Yes	Yes	No	No
Industry $\times$ province $\times$ year FE	No	No	No	No	No	No	Yes	Yes
Observations	1,870,095	1,869,370	1,819,795	1,819,090	1,819,795	1,819,090	1,805,305	1,804,600
R^2	0.345	0.345	0.345	0.345	0.346	0.346	0.359	0.359

Notes: This table reports estimates for the population of married entrepreneurs in Canada. The dependent variable is firm investment, measured by asset growth. The main explanatory variable is the change in the layoff rate in the industry in which the entrepreneur's spouse is employed. Controls include polynomials in firm age and in the entrepreneur's age. Even-numbered columns additionally control for whether the spouse was laid off and for the spouse's log earnings. 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.8: Heterogeneous effects on investment by firm age

	Young firms				Older firms			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post × $\text{Oil}_{i(j)}$	-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 × 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	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. $\text{Oil}_{i(j)}$ 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. 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.9: Robustness to finer geographic effects: postal codes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post $\times$ $\text{Oil}_{i(j)}$	-0.126**	-0.126**	-0.142**	-0.131*	-0.0717***	-0.0564**	-0.0772	-0.0863*
	(0.062)	(0.062)	(0.072)	(0.069)	(0.0207)	(0.0202)	(0.0438)	(0.0438)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry $\times$ year FE	No	No	Yes	Yes	Yes	Yes	No	No
Postal code $\times$ year FE	Yes	Yes	Yes	Yes	No	No	Yes	Yes
Observations	3,635	3,635	3,000	3,000	11,430	11,430	2,775	2,775
R^2	0.658	0.668	0.706	0.716	0.347	0.369	0.405	0.428

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. The dependent variable 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. $\text{Oil}_{i(j)}$ 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. Columns (1)–(4) use the baseline matched sample. In Columns (5)–(8), the sample is re-constructed by re-matching treated and control observations using postal code as the geographic matching variable. Columns (5)–(6) match observations using postal code only. Columns (7)–(8) use the same matching criteria as in the main specification, except that commuting zone is replaced by postal code and industry is measured at the 1-digit level. 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.10: 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 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 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 describes the algorithms used to solve the model. The solution nests an outer bisection on r (Algorithm 6), a direct wage computation from no-arbitrage (Algorithm 1), value-function iteration for singles and couples (Algorithms 2–3), and forward iteration for stationary distributions (Algorithm 4).

B.0.1 Grid Specification

Wealth a is placed on a power-spaced grid of N_A points spanning $[0, a_{\max}]$,

$$a_i = a_{\min} + (a_{\max} - a_{\min}) \cdot \left(\frac{i}{N_A}\right)^p, \qquad i = 1, \ldots, N_A, \tag{16}$$

with $p = 2.5$ providing denser coverage near the low-wealth region where policies change rapidly. Entrepreneurial productivity z_e follows a 5-state Tauchen discretisation of the AR(1) log process ($\rho_z = 0.90$, $\sigma_z = 0.20$), with levels normalised by their ergodic mean so that $\sum_j \pi_j z_{e,j} = 1$. Worker productivity $z_w \in \{z_w^L, z_w^H\}$ is binary and permanent. Employment state $ie \in \{0, 1, 2\}$ is discrete.

B.0.2 Wage and Corporate Capital from No-Arbitrage

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. $\text{KN} \leftarrow \left[(r + \delta_k)/(\alpha\, A)\right]^{1/(\alpha-1)}$
2. $w \leftarrow (1 - \alpha)\, A\, \text{KN}^{\alpha}$
3. After solving the agent problem: $N_C \leftarrow \sum_{i \in \text{employed}} z_{w,i}$; $K_C \leftarrow \text{KN} \times N_C$.
4. **return** w, KN, K_C

B.0.3 Singles Value Function Iteration

Algorithm 2: Singles VFI (for a given z_w)

Data: Rate r, wage w, TFP A, permanent type z_w, parameters

Result: Value function $V(a, z_e, ie)$, policies, occupational choice

1. Discretise state space: $\{a_i\}$ power-spaced, $\{z_{e,j}, P_z\}$ Tauchen, $ie \in \{0, 1, 2\}$.
2. Initialise $V \leftarrow 0.05 \cdot a$.

Main Loop For $n = 1, 2, \ldots$:

Compute CE arrays: $\text{CE}_E(a')$ (survival mixture), $\text{CE}_{W,0}(a')$ (Markov σ_s), $\text{CE}_{W,1}(a')$ (Markov φ_f).

For each (a, z_e, ie): evaluate V_E and V_W, take the max.
Howard policy evaluation: H steps holding policies fixed.
Converge when $\|V^{(n+1)} - V^{(n)}\|_\infty < \text{TOL}_{\text{VFI}}$ or policy-stable for 8 iterations.

3. **return** V, policies.

B.0.4 Couples Value Function Iteration

Algorithm 3: Couples VFI (for a given $(z_{w,1}, z_{w,2})$)

Data: Rate r, wage w, TFP A, types $(z_{w,1}, z_{w,2})$, parameters
Result: Value function $V(a, z_{e,1}, z_{e,2}, ie_1, ie_2)$, policies, config choice

1. Discretise state space $N_A \times N_Z^2 \times 3^2$.
2. Precompute EW and WE profit arrays with separate collateral terms.

Main Loop For $n = 1, 2, \ldots$:
Compute four CE arrays (EE, EW, WE, WW) with config-specific employment mixtures.
For each state, evaluate all four configs and take the max.
Convergence as in singles VFI.

3. **return** V, policies, config choices.

Called four times per excess-demand evaluation: once for each $(z_{w,1}, z_{w,2}) \in \{z_w^L, z_w^H\}^2$.

B.0.5 Stationary Distribution

Algorithm 4: Stationary Distribution

Data: Policies from VFI, grids, transition matrices
Result: Stationary distributions; aggregates

1. Initialise $\mu^{(0)}$ uniformly on assets, stationary on (z_e, ie).
2. Forward-iterate for N_{dist} steps, distributing mass via linear interpolation.
3. **return** stationary distributions and aggregates (K_E, N_C, etc.).

B.0.6 GE Solver

Algorithm 5: Excess-Demand Evaluation

Data: Candidate rate r, parameters, TFP A, warm-start V_0
Result: ED(r), all equilibrium objects

Step 1: $(w, \text{KN}) \leftarrow$ **Algorithm 1**(r, A).

Step 2: Solve 6 VFI problems: 2 singles (Alg. 2) + 4 couples (Alg. 3).

Step 3: Forward-iterate stationary distributions for each type.

Step 4: Aggregate: weight singles by $\Pr(z_w)$, couples by matching probs.

Step 5: $\text{ED}(r) = (K_C + K_E) - A_{\text{tot}}$.

Algorithm 6: GE Solver: Bisection on r

Data: Parameters, TFP A, optional r_{hint}, warm-start V_0

Result: Equilibrium r^* and all objects

1. Bracket $[r_{\text{lo}}, r_{\text{hi}}]$ around r_{hint}.
2. Bisect until $|\text{ED}| < \text{TOL}_{\text{GE}}$ or $r_{\text{hi}} - r_{\text{lo}} < \text{TOL}_r$.
3. **return** r^* and all equilibrium objects.

B.0.7 Computational Parameters

Table B.11: Computational Parameters

Parameter	Symbol	Value	Description
Asset grid points	N_A	120	Power-spaced, $[0, 60]$
Productivity states	N_Z	5	Tauchen approximation
Employment states		3	$ie \in \{0, 1, 2\}$
Grid exponent	p	2.5	Power spacing
VFI tolerance	TOL_{VFI}	1×10^{-5}	Sup-norm on V
Max VFI iterations		1000	
Howard policy-eval steps	H	20	Per VFI iter
GE tolerance	TOL_{GE}	5×10^{-3}	$\lvert\text{ED}(r^*)\rvert$
Max GE iterations		50	Bisection on $\text{ED}(r)$
Interest rate bracket	$[r_{\min}, r_{\max}]$	$[-0.02, 0.08]$	
Bracket-collapse tol.	TOL_r	10^{-5}	Early exit
Distribution iterations	N_{dist}	6000	Forward iteration