Monetary Shocks and Labor Income Risk with a Billion Observations*

Mario Giarda† Ignacio Rojas‡ Sergio Salgado§

March 21, 2025

Preliminary and Incomplete

Abstract

We investigate the impact of monetary policy shocks on labor income risk. Using high-quality administrative data from Chile, we show that after a monetary shock, the skewness of the earnings growth distribution declines, with no significant response of the dispersion of earnings growth. In contrast to what is observed in a typical recession, the skewness drops because positive earnings changes become less likely after the monetary shock, but not because wage cuts increase in probability. The drop in skewness is most prominent for high-earnings workers and for workers who stay with the same employer. Using a standard heterogeneous agents incomplete markets model, we show that the shift in income risk is an important mechanism for the transmission of aggregate shocks.

*Preliminary and Incomplete. The views expressed are those of the authors and do not necessarily represent the views of the Central Bank of Chile or its board members. This study was developed within the scope of the research agenda conducted by the Central Bank of Chile (CBC) in the economic and financial affairs of its competence. The CBC can access anonymized information from various public and private entities due to collaboration agreements signed with these institutions.

†Central Bank of Chile; mgiarda@bcentral.cl

‡Central Bank of Chile; irojas@bcentral.cl

§The Wharton School, University of Pennsylvania; ssalgado@wharton.upenn.edu

1 Introduction

A significant body of evidence indicates that workers' risk varies over the business cycle. For instance, Storesletten et al. (2004) shows that the variance of income shocks is significantly higher during recessions. More recently, empirical evidence based on administrative data has shown recessions are characterized by an increase in left-tail risk, as measured by the procyclicality of the skewness of earnings growth (see for instance, Guvenen et al., 2014 and Guvenen et al., 2022 for international comparison). As such, changes in risk add to the welfare cost of the business cycle. On the other hand, a growing body of literature studies the impact of aggregate shocks–and policy changes– on worker heterogeneity and how this heterogeneity affects the transmission of these shocks to macroeconomic aggregates. This literature, however, typically assumes that income risk follows an exogenous–and static–process that is independent of policy changes (See Kaplan et al., 2018 or Auclert et al., 2021 among others).

In this paper, we study income risk over the business cycle; we study how the cross-sectional moments of the income growth distribution respond to monetary policy shocks. In our analysis, we use high-quality administrative data from Chile that contains monthly-level information on labor earnings, a firm identifier, as well as other worker characteristics such as gender, age, and education. Our baseline sample contains about a billion worker-month observations covering 17 years of data from 2005m1 to 2021m12. Using this information, we construct measures of earnings growth at different horizons, from one to twenty-four months, and corresponding employment transitions. Then, we compute the cross-sectional moments of earnings growth and study their responses to monetary policy, which allows us to evaluate how monetary policy affects income risk at short and long horizons.

In our main empirical analysis, we run a series of Local Projections (Jordà, 2005) on different cross-sectional moments of labor earnings growth distribution. We use the shocks built by Aruoba et al. (2021) for Chile, who consider as monetary surprises the difference between the actual change in the monetary policy rate and expectations by forecasters the day of the meeting before it takes place. We focus on the impact of monetary shocks on the average, dispersion, skewness, and kurtosis of the cross-sectional distribution of the 12-month changes in workers' labor earnings. As expected, after a contractionary monetary shock, average earnings growth declines. Dispersion, however, does not respond significantly to changes in monetary policy. Instead, it is the skewness of earnings growth that is more responsive to a monetary shock: after an unexpected increase of

the monetary policy rate by 25 basis points (bp.), the Kelley skewness of annual earnings growth declines by 3 percentage points corresponding to 15% of the fall observed during the Global Financial Crisis. Hence, monetary policy contributes significantly to fluctuations in labor income risk. Interestingly, it is the right tail of the distribution of earnings growth that is more responsive to a monetary shock, indicating that monetary policy reduces the likelihood of earnings gains rather than increasing the probability of earnings cuts. The kurtosis, instead, tends to increase after a monetary shock, indicating that a monetary shock makes extreme changes in earnings less likely.

We also find significant heterogeneity across workers of different earnings levels and for workers who experience different employment transitions. In particular, high-income workers tend to experience a significantly larger increase in income risk after a monetary shock relative to low-income workers. This increase in risk translates into a higher dispersion of the income growth distribution, but more importantly, into a decline in the skewness.

Additionally, skewness declines for workers who maintain a stable employment status (not experiencing unemployment transitions). To show this, we study the evolution of income growth moments for different employment state transitions between period t and $t+h$: transitions to and from unemployment; job-to-job transitions; and the risk faced by workers who stay in their firms, that we refer as stayers. We find that the response of the skewness of earnings growth for stayers and job-to-job transitions responds significantly more to a monetary shock than the skewness of earnings growth for workers who get unemployed. Therefore, unemployment is not the largest contributor to the earnings risk generated by a monetary policy shock. This result implies that income risk matters beyond transitions between employment states.

Finally, we evaluate quantitatively whether and by how much income risk faced by all agents affects aggregate consumption. We study an otherwise standard heterogeneous agent, incomplete markets model to do that. Relative to the standard model, here individuals face time-varying risk, in the form of a skewness shock to earning growth, which responds to aggregate shocks, which presumably can be driven by policy decisions, such as monetary policy. Our results suggest that variation in earnings growth is a quantitatively relevant contributor to fluctuations in aggregate consumption.

This paper is related to McKay (2017a), who models income risk as the probability of having long-run unemployment; in our case, we want to consider pure income risk—understood as the income risk suffered by people who remain in their jobs and not employment flows' risk. This paper also relates to Bayer et al. (2019). Nonetheless, we want to consider the third moment of

the income change distribution as well as capture all the nonlinearities of this problem since they focus on the effect of the second moment in a second-order approximation of the model. Also, this paper is related to Debortoli and Galí (2024), who study the role of idiosyncratic risk in aggregate consumption and find that if rich consumers face a higher uncertainty, idiosyncratic risk has an amplification effect. We show empirically that this is the case and extend that idea to fluctuations in income risk.

The rest of the paper is organized as follows. In Section 2 we describe our data and the main characteristics of the distribution of labor earnings growth at different horizons. In Section 3 we perform our main empirical exercise that relates monetary shocks and labor income risk whereas Section 4 extends our main results for workers of different characteristics. Section 5 discusses our quantitative exercise. Finally, section 6 concludes.

2 Data and Stylized Facts

In this section, we describe our data construction based on Chilean administrative data, and then we present some stylized facts on the distribution of earnings growth over the business cycle.

2.1 Data

We draw information from the administrative dataset of the unemployment benefit management institution (AFC, by its Spanish acronym) to obtain the time series for wages and labor transitions for the universe of employee workers in Chile. The raw labor data contains information about workers' monthly salaries, the use of unemployment benefits, and employer-employee matches from Jan-2002 to Dec-2024. This data also includes the educational level, gender, firm industry, and amount saved in the unemployment insurance account, among other variables. The database covers all formal workers in the private sector (workers with a formal contract who pay pension contributions). The total number of observations is about 900 million, with monthly observations hovering around six million per month since 2014. To account for different transitions between jobs and in and out of the sample (Job-to-Unemployment, for example), we fill with zeros the months without observed wages as explained below, which allows us to reach about a billion observations in total for the sample from Jan-2005 to Dec-2021.[1]

We complement this dataset with aggregate variables such as monthly data on industrial

[1]See Appendix A for additional details of our database.

Table 1: DESCRIPTIVE STATISTICS FOR BASELINE LABOR EARNINGS SAMPLE

Sample Statistics, Employed							
	Total number of		Ratio of	Total number of			Total
Year	Workers	Firms	Men	$w_{it} > 0$	$w_{it} = 0$	$w_{it} < \underline{w}_t$	Observations
All ages workers	8,912,786	1,452,258	56.34%	715,757,539	392,068,363	576,198,553	1,107,825,902
[22,60] age workers	8,407,093	1,419,838	56.64%	671,622,988	340,080,526	463,434,849	1,011,703,514

Note: This table show summary statistics of our database covering the period 2005m1 to 2021m12.

production (IMACEC), CPI, exchange rates, unemployment, real wages, terms of trade, and the monetary policy rate (MPR). We extract these data from the Central Bank of Chile's Statistical Database.

Sample Selection. Because of the high quality of our dataset (which contains information directly reported by firms, with little to no misreporting and no sample attrition), we do not conduct a significant cleaning of our database to have the largest sample possible each period. At least for the twelve-month earnings changes, which is the main focus of our analysis, we are less concerned about the instabilities that people with short spells in the labor market may induce in our result and potential seasonal effects. We are considering year-on-year changes in the time-series analysis below, which requires people to be in the sample for two years. The only cleaning we do is with people with two jobs and on sick leave, where we consider the largest income in the former case and impute the last available income in the latter. Also, our database is top-coded, with about three percent of the database at the top. This is not a limitation for our exercise since we would likely underestimate income changes.

As we explain below, in our sample, we follow workers who find a (formal) job at least once. We fill those periods with a zero if they do not appear in our sample after entering. This is valuable information for building the tails of the income growth distribution with transitions between jobs and unemployment. We label the zero state as unemployment (recognizing that "unemployed" workers can be employed in the informal or public sectors) and impute a wage. The number of observations we have with positive income in the sample 2005-2021 is 715,757,539, and observations with zero income, which we impute, are 392,068,363. Thus, we reach a massive dataset of more than a billion monthly observations. These numbers are after the cleaning we make. We have about 9 million workers' IDs and 1.4 million firm IDs. The percentage of men is 56.34%.

Definition of Earnings. The main measure of labor earnings considers all sources of labor income, including base pay, tips, and bonuses. We do not include any measure of self-employment income. We denote these earnings by w_{it}. Then, we take logs, and for those observations below the minimum income ($\underline{w}_t$), we follow the procedure described in Guvenen et al. (2015), and transform income as follows:

$$y_{it} = \begin{cases} \log\left(\underline{w}_t + \frac{\underline{w}_t}{10(\underline{w}_t+10)}\left\{(w_{it}) + 10U[0,1]\right\}\right) & \text{if } \; w_{it} \leq \underline{w}_t \\ \log w_{it} & \text{if } \; w_{it} > \underline{w}_t. \end{cases}$$

This way to account for monthly incomes below the real minimum income has the following advantages. It preserves the order of the income distribution and allows for capturing moments related to flows, including unemployed individuals whose responses across the business cycle are relevant, such as the UtJ and JtU flows. Finally, fluctuations with zeros represent compositional effects, meaning that higher skewness from separation must be due to higher-income workers losing jobs. Moreover, relative to the standard arc-percentage change, this measure is unbounded and does not generate large concentration of earnings observations in -2 or +2.[2]

Growth Rates and Moments of the Distribution. Define income growth between h periods as $\Delta^h y_{it} = y_{it} - y_{it-h}$. With this, we calculate (and study) the moments of the cross-sectional distribution of $\Delta^h y_{it}$. The moments we consider are the mean, variance (dispersion), skewness (asymmetry), and kurtosis (tailedness). To build high-order moments, we use their quantile-based measures (following Guvenen et al., 2014). This is, we define dispersion as the $P90 - P10$ difference. We use the Kelley skewness,

$$\text{Skewness} = \frac{(P90 - P50) - (P50 - P10)}{P90 - P10},$$

as our measure of asymmetry. And finally, we use the Crow-Siddiqui measure,

$$\text{Kurtosis} = \frac{P97.5 - P2.5}{P75 - P25} - 2.91,$$

[2]The arc-percent is defined as $d = \frac{x_t - x_{t-1}}{0.5 \times (x_t + x_{t-1})}$ where one of the x can be equal to 0. This measure is bounded between -2 and 2.

Table 2: Time series means and correlations with output.

Moment	1-M	6-M	12-M	24-M
Avg_t^h	0.01	0.03	0.06	0.11
$Disp_t^h$	0.46	1.18	1.33	1.65
$Skew_t^h$	0.07	0.07	0.08	0.08
$Kurt_t^h$	134.5	14.03	10.72	7.01
$corr(Avg_t^h, \Delta^h Y_t)$	0.31	0.62	0.77	0.69
$corr(Disp_t^h, \Delta^h Y_t)$	-0.09	0.32	-0.21	-0.59
$corr(Skew_t^h, \Delta^h Y_t)$	0.43	0.76	0.85	0.81
$corr(Kurt_t^h, \Delta^h Y_t)$	-0.03	-0.09	-0.57	-0.28

Note: This table shows time series average cross-sectional moments of the distribution of earnings growth for 1- to 24-months changes. Dispersion is measured by the P90-P10 percentiles differential, skewness by the Kelley Skewness, and kurtosis by the Crow-Siddiqui measure. The bottom panel shows the correlation of each moment and the corresponding change in GDP as approximated by the change in industrial production index (IMACEC).

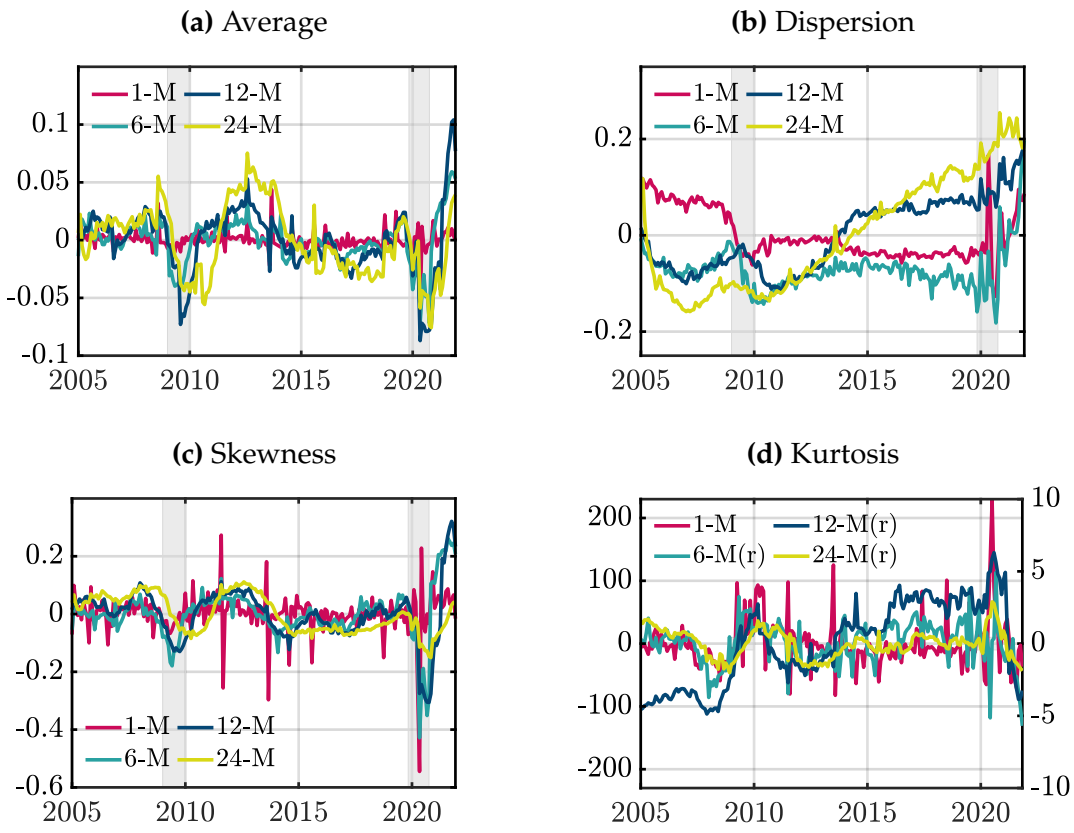
as our measure of tailedness. These measures are more robust to outliers and offer a convenient way to decompose them. Next, we show the evolution of these moments over the business cycle.

2.2 Characterizing the Distribution of Earnings Growth

The top panel of Table 2 shows sample means of the first four moments of the distribution of earnings growth at four different horizons: 1, 6, 12, and 24 months. Over the period we analyze, Chile experienced consistent earning growth, captured by a sample average of the average growth that increases as we move to longer horizons. Dispersion increases as well as we move from 1 to 24 months, as workers experience job mobility and changes in their income over time. Consistent with an overall increase in real earnings, the skewness of the distribution—measured by the Kelley measure—is positive along all the horizons and more or less of the same magnitude. As for the kurtosis, this tends to decrease as we move from shorter to longer horizons.

The bottom panel of Table 2 shows the correlation between each moment and GDP growth at the corresponding horizons (measured by the industrial production index, IMACEC). As expected, the average income growth is positively correlated at all horizons, whereas the skewness is positively correlated, as it is typically found in the literature. The dispersion, however, does not show a clear pattern. Figure 1 depicts this more clearly, which shows the time series of the first four moments of the distribution of earnings growth with the shaded areas indicating recessions. Visually, Figure 1 shows the same as above: the average and the skewness fall in recessions for all horizons, whereas

Figure 1: Moments of the Distribution of Monthly Earnings



Note: This figure shows the time-series of the moments of the cross-sectional distribution of income growth at different horizons. We show the average, dispersion, skewness and kurtosis. Each plot shows a moment. Each figure shows the moment of 1-, 6-, 12-, and 24-month changes distribution. The figure considers data from 2005m1 to 2021m12. The data is detrended and deseasonalized (when it applies) with ARIMAX-13. Shaded areas represent recessions defined as periods with more than 6 consecutive months of y-o-y drop in GDP.

the dispersion does not seem to have a systematic behavior over time. Kurtosis, instead, tends to increase in recessions. These patterns—based on monthly data—are consistent with the existing literature (Guvenen et al., 2014) which typically focuses on annual changes in annual earnings. In the next section, we study the behavior of these moments over the business cycle more in-depth how they respond to a monetary policy shock.

3 Income Risk Responses to Monetary Shocks

We identify the impact of monetary policy shocks on labor income risk using a series of Local Projections (Jordà, 2005) of the following form,

$$Y_{t+k} - Y_{t-1} = c + \beta^k mps_t + \Gamma' X_t + \delta_t + \varepsilon_{kt}, \tag{1}$$

where $Y_{t+k} - Y_{t-1}$ is the cumulative response in the variable of interest (e.g., the dispersion of earnings growth), between months t and $t+k$. The variable mps_t corresponds to a monetary policy shock, defined as the difference between the Monetary Policy Rate (MPR) defined at the Chilean Central Bank Board meeting on a particular date and the expected rate provided by analysts as collected by Bloomberg the day of the meeting (Aruoba et al., 2021).[3] We show the time series of the shock in Appendix Figure B.1. The control vector X_t includes the level of the MPR, twelve-month CPI (consumer price index) inflation, the 1-month log difference of monthly output, the 1-month log change of the Terms of Trade index, and the 1-month log change of the nominal exchange rate. The term δ_t is a linear trend. The choice of these controls is similar to Aruoba et al. (2021); however, we make two additional assumptions to obtain a proper response of the CPI to our monetary shock. First, instead of including a log of CPI in the controls, we include twelve-month inflation. The reason is that in Chile, the inflation rate is extremely stable over the period we studied. Thus, we need a longer horizon in the lags of the CPI. The twelve-month change in CPI captures these longer lags. Second, to avoid the well-known price puzzle (see, e.g., Christiano et al., 1999), we control for terms of trade, which is an important driver of inflation in a small open economy like Chile.[4]

We use a twelve-lag of our control variables and consider the sample between September 2005 and December 2019 in our Local Projections exercise. Data after December 2019 is available but heavily affected by the COVID-19 pandemic. When we study the response of the moments of the income growth distribution, for instance, the skewness of earnings growth, we also control for lagged moments' change of the moment. Finally, we present the smooth local projections, following Barnichon and Brownlees (2019), which smooths out the Local Projections using a ridge regression of impulse response smoother parameters. We are interested in the sequence of parameters $\{\beta^k\}_{k=0}^K$, which represents the change in variable Y_t between t and $t+k$ due to the monetary policy surprise. Also, in all plots, we show the 90 percent confidence bands.

[3] Since 2019 the board of the Central Bank of Chile meets 8 times a year in fixed months.

[4] MPR, CPI, output (IMACEC), Terms-of-Trade, and nominal exchange rate were all taken from the Central Bank of Chile's statistical database.

3.1 Responses of Aggregate Variables

We start our analysis by showing the response of aggregate variables to a monetary policy shock. Figure 2 shows the response of GDP, unemployment, CPI, and the real wage to a monetary policy surprise of 25bp.[5] All the aggregate responses align with other studies for Chile (Aruoba et al., 2021) or the US (Jarociński and Karadi, 2020). After a contractionary surprise, output (Figure 2a) decreases significantly after 3 to 4 months, with a bottom response at 18 months of 1.2 percent. The economy remains below trend for about 24 months after the shock. As Figure 2b shows, unemployment increases significantly, about 0.2 percentage points 4 months after the shock, increasing until it reaches a 0.6 p.p. peak at 18 months. Notice that the timings between output and unemployment responses are very similar, suggesting a strong cyclicality of the labor market in Chile, also conditional on monetary policy.

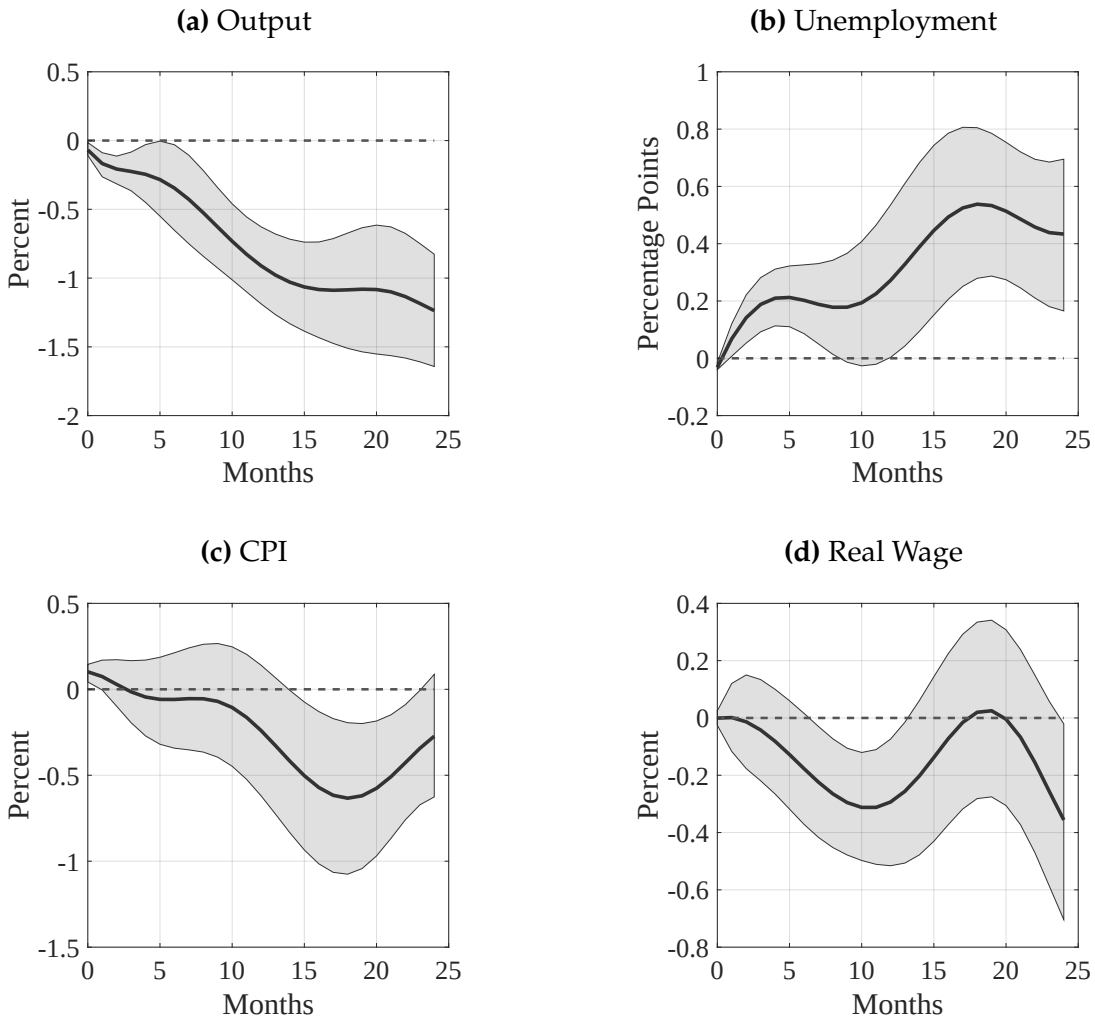
Figure 2c depicts the responses of CPI and Figure 2d real wages to the monetary surprise. We find that for our specification, prices fall—significantly—in response to monetary surprises twelve months after the shock. The CPI response is milder than the output response with the trough at about 18 months after the shock. With this result, we are confident that our specification delivers a proper response of the aggregates to a monetary policy shock, generating a fall in output and CPI. Finally, the real wage falls after the monetary policy shock. The response of wages is slower than that of real variables, being significant after nine months. The real wage recovers at the time of the trough in the CPI, but in the long run, it becomes highly negative.

3.2 The response of the earnings distribution to a monetary shock

The next step is to study how monetary shocks affect the distribution of income growth. For that, we run the specification in Equation (1), but we replace the moments of the monthly income growth distribution on the left-hand side. This exercise only considers the moments of the distribution of earnings growth over a twelve-month horizon, $\Delta^{12} y_{it}$, including the same controls and assumptions discussed above for the aggregate variables. We use annual changes of monthly earnings for three reasons. First, this considers any potential seasonality in earnings data, which is quite high at monthly frequencies. Second, considering twelve-month earnings changes allows for a more direct comparison with the existing literature that typically focuses on year-to-year changes in annual earnings (Guvenen et al., 2022). Third, in the following sections, we study how the distribution of

[5] Unemployment and real wages are obtained from the Central Bank of Chile's statistical database.

Figure 2: Responses of aggregate variables to monetary surprises.



Note: This figure shows the response of macroeconomic variables to a monetary surprise of +25bp. The figure considers data from 2005m9 to 2019m12. Output is the Industrial Production Index (IMACEC). In each Figure, we show the smoothed local projection estimates. Shaded areas are 90% confidence intervals.

earnings growth changes for workers experiencing different labor market transitions (e.g., Stayers, Job-to-Job, and so on), for which one would prefer to have a few months to observe better workers experiencing different labor market transitions. Fourth, year-to-year changes are still between the business cycle frequency based on expected income changes. For completeness, however, we report the response of the distribution of earnings growth of one-month changes in Appendix D. We find qualitatively similar results to the ones we provide below with 12-month changes.

Figure 3 depicts the response of four moments of the income growth distribution to monetary shocks. We consider the average, dispersion (P90-P10 differential), skewness (Kelley measure),

and kurtosis (Cross-Siddiqui measure). We find that the average drops initially by about 0.5% by 4 months after the shock, while we observe a trough of about 1% by twenty months after the monetary shock. This is, income grows one percent less twenty months after a monetary shock, recovering starting levels about two years after the shock. These are strong effects, considering that the shock size is only 0.25%. Second, we find little impact from monetary policy on the dispersion of income growth, which has an oscillating response that we interpret as zero effect. This last finding is consistent with the unconditional evidence presented above, and it confirms the findings in previous literature (e.g., Guvenen et al., 2014) that dispersion of earnings growth does not significantly fluctuate over the business cycle.

Figure 3c shows the response of the skewness to the monetary surprise. As a result of a monetary shock, the skewness falls significantly after one month. Left tail risk in the Chilean labor market rises quickly, reaching 0.02 points of lower skewness six months after the shock. The trough of the skewness is at -0.035 points 18 months after the shock. Then, the skewness recovers, reaching pre-shock levels 28 months after the shock. We consider these effects economically significant. Figure 1 shows skewness fell by 0.28 during the Great Recession. Since the response to a 0.25% monetary surprise is -0.035, a small monetary shock corresponds to 15 percent of the fall in the skewness during a very tough recession. If we extrapolate this linearly, a monetary shock of 1% would generate about 60 percent of the fall in the skewness we observed during the Great Recession.

Figure 3d shows the response of kurtosis of earnings growth to a monetary surprise. Recall that the kurtosis measures how much of the dispersion of a distribution is on the tails relative to the center, with a higher number indicating a more significant importance of the tails. In response to a monetary shock, kurtosis increases significantly; that is, the mass of extreme values relative to the central mass becomes larger in response to monetary policy shocks. Excess kurtosis increases by about 0.8 relative to the trend 20 months after the shock, which makes the earnings change distribution more leptokurtic. Relative to the mean of 13 over our sample, the kurtosis increases significantly; taken to the change after a 1% for instance, an increase in the monetary policy rate, implies an increase of 3.2 in the kurtosis. We conclude that a contractionary monetary policy has a significant impact on the distribution of earnings growth, and raises labor income (left-tail) risk without a significant increase in overall risk (as measured by the dispersion).

3.3 Decomposing the response of higher-order moments.

To understand the sources of income risk further, we study the responses of the components of our measures of skewness and kurtosis. Recall that the Kelley skewness is given by $((P90 - P50) - (P50 - P10))/(P90 - P10)$ so we can study the responses of $(P90 - P50)$ and $(P50 - P10)$ separately to evaluate which tail accounts for the decline in skewness. Figure 4 shows the responses of these components to the monetary surprise. To fix ideas, recall that movements in $(P90 - P50)$ are changes in the right tail, meaning that large-positive income changes are more or less likely, whereas changes in $(P50 - P10)$ are left tail changes, which means that large negative changes are more or less likely.

In response to the monetary surprise, Figure 4a shows that $(P50 - P10)$ increases, meaning that monetary contractions make large negative changes more likely. On the other hand, Figure 4b shows that $(P90 - P50)$ decreases after the monetary policy shock, implying that the right tail becomes narrower. We find that the response of the right tail is much more pronounced than that of the left tail. This suggests that monetary policy works by driving a loss of income for workers who would have had significant raises in earnings in the absence of the shock rather than by generating large earning cuts, as is typically observed during a recession, where the P50-P10 differential explains most of the decline in the skewness of the distribution (Guvenen et al., 2014). In fact, in Appendix C, we show that during the two recessions present in our dataset, it is the P50-P10 differential that changes the most. So, the relative importance of right- versus left-tail is specific to a monetary policy shock.

Figure 5 shows the decomposition of the Crow-Siddiqui kurtosis. Recall that this measure is given by $(P97.5 - P2.5)/(P75 - P25)$ where the numerator $(P97.5 - P2.5)$ gives the mass at the extremes, and the denominator gives the mass at the middle of the distribution. This is a measure of the ratio of the extremes to the middle of the distribution and gives a measure of the "tailedness" of the distribution. Figure 5a shows that the numerator increases in response to the shock, while Figure 5a shows that the denominator goes down. Thus, extreme values become more likely. An alternative way to interpret this result is that changes around the mean are more likely than those between the mean and the extremes of the distribution.

4 Monetary policy shocks, income risk, and worker characteristics

In this section, we study how the response of the income distribution to a monetary stimulus varies across workers of different characteristics. We focus on two main exercises. First, we analyze how the response of the different moments of the earnings growth distribution varies across workers of different lifetime earnings—a measure of their fixed type. Second, we analyze how the response of different moments of the earnings growth distribution varies depending on labor-market transitions: employment to unemployment, employment to employment, and so on—a measure of workers' exposure to the business cycle.

4.1 The response of earnings growth across workers types

To understand how workers of different characteristics respond to monetary policy, we need to separate the role of firms in this analysis.[6] To keep the exercise as simple as possible, we estimate a two-sided fixed effect regression in the spirit of Abowd et al. (1999) (AKM).

More precisely, we follow the procedure of Card et al. (2018) and use the panel of earnings levels to estimate the following model linear wage model

$$\log Y_{it} = \phi_i + \varphi_{F(i,t)} + \psi_{Age(i,t)} + \nu_{(Year_t)} + x'_{i,t}\beta + \varepsilon_{i,t},$$

where ϕ_i is the individual fixed effect, $\varphi_{F(i,t)}$, $\psi_{Age(i,t)}$, and $\nu_{(Year_t)}$ are the firms, age group, and year fixed effects, respectively. The ϕ_i measures a worker's permanent income. Then, we can classify workers into quintiles of the fixed effects distribution (our measure of worker permanent earnings) and calculate the moments of the cross-sectional distribution for each quintile. In our main exercise, we study the responses of these (conditional) moments to a monetary policy using the same local projections described in Section 2.[7]

Figure 6 shows the responses of the moments of the income change distribution to a monetary surprise by quintile of permanent earnings. We show the average, dispersion, skewness, and kurtosis. We show the aggregate responses and also the responses of the different quintiles. To better distinguish the responses of different groups, we display the top quintiles in blue and the bottom quintiles in red.

[6] It is well known that firms pass a significant fraction of their shocks to workers' wages (Guiso et al. (2005) and Guiso and Pistaferri (2020)) and that this passthrough is heterogeneous, with larger firms having a lower passthrough (Chan et al. (2019)). Although interesting, our focus here is on workers' characteristics.

[7] A few papers have implemented AKM-style regressions using Chilean data. See for instance Aldunate et al. (2020).

Figure 6a shows the response of averages. The averages of the twelve-month distribution respond significantly heterogeneously, with a stronger response of bottom quintiles, except for the first quintile. This result is because workers at the bottom of the earnings distribution typically earn the minimum income, which does not significantly respond to a monetary policy shock.[8] Similarly to what we find in the aggregate response of average earnings—reproduced in the plot in black for comparison–the dispersion response is close to zero except for the extreme quintiles: Q1 dispersion falls while Q5 dispersion increases. However, the response of the Q5 is significantly higher than Q1. In other words, earnings risk in response to a monetary policy shock increases the most for high-earnings workers.

The main result of this section is the response of skewness by quintile. The skewness for the bottom two quintiles drops by less than the three top quintiles. This implies that those suffering the most from earnings risk fluctuations in response to a monetary policy shock are workers from the third to the fifth quintiles, which explains the aggregate drop in skewness. In other words, although a contractionary monetary policy shock tends to reduce earnings growth for most workers, those workers at the top mostly experience an increase in the probability of significant adverse changes in earnings, which is captured by a skewness that becomes more negative. Finally, we find little differences in the responses of the kurtosis between quintiles except for the fifth quintile.

4.2 Labor Market Flows

Several papers have documented the impact of monetary policy shocks on workers' labor market transitions (see, e.g., Hubert and Savignac, 2024 and Gulyas et al., 2024). Furthermore, empirical evidence shows that the distribution of earnings growth for workers that change employers is more dispersed and more left-skewed relative to workers that maintain stable employment relations (Guvenen et al., 2021). In this section, we link these two strands of literature and study the role of labor market transitions in the response of the skewness and kurtosis to monetary surprises.

To do so, we build different labor market transitions over the period we analyze income growth. That is, for a given growth rate $y_{it+h} - y_{it}$, we track the transitions (or not) between t and $t+h$, as follows: i) Stayers: individuals who work in the same firm in t and $t+h$. ii) Job-to-Job (JtJ): individuals who work in different firms in t and $t+h$. iii) Job-to-Unemployment (JtU): individuals who work in t and are unemployed in $t+h$. iv) Unemployment-to-Job (UtJ): individuals unemployed in t and with a job in $t+h$.

[8] For evidence on this fact for income levels in Chile, see Aldunate et al. (2025).

Table 3: Moments of the Distribution of Earnings Growth for Different Transitions (2015).

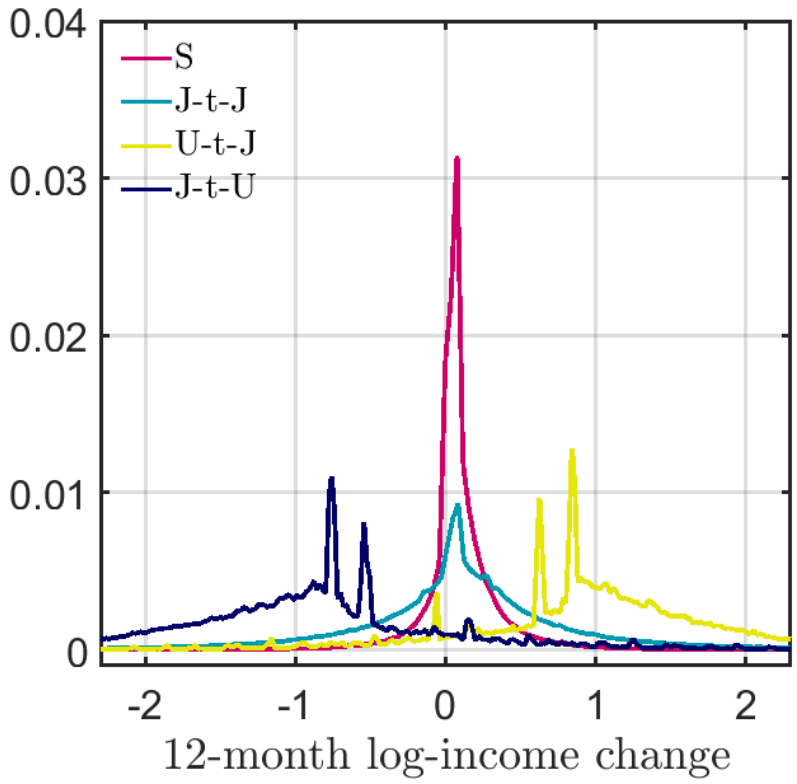
Moment	Stayers	JtU	JtJ	UtJ	Total
Average	0.07	-0.79	0.11	0.85	0.07
Dispersion	0.51	2.07	1.70	2.00	0.6
Skewness	0.19	0.02	0.04	0.00	0.10
Kurtosis	5.38	1.96	2.62	1.98	7.59

Note: This table shows moments of the distribution of earnings growth within workers who experience different transitions in the labor market over 12 months. We show the moments of 2015.

Table 3 describes the distribution of earnings growth after each of the transitions discussed before. Most of the observations in our sample correspond to stayers. We have that 39 percent of the sample is growth rates of workers who remained in the same firm. Many observations in our sample are people who were unemployed in both periods, which is about 24 percent of the sample. Then, JtJ transitions comprise a large proportion of our sample, 16 percent. Finally, JtU and UtJ correspond to about 10 percent each. In our sample, the remainder of transitions correspond to a tiny proportion.

Table 3 and Figure 7 characterize the distribution of earnings growth for each transition. Income growth is mainly due to JtJ transitions, with an average yearly growth of 11 percent. During this period, there was a significant increase in labor earnings in Chile, generating an average of 7 percent annual earnings growth for stayers. These two distributions are right-skewed. Naturally, transitions from unemployment and to unemployment imply positive and negative growth, respectively, but they by themselves are also skewed. JtU is negatively skewed, while the skewness of workers who experience a UtJ transition is mildly negative. This suggests significant income losses from unemployment. Figure 7 shows how these flows are distributed along income changes, with Stayers and JtJ centered in a mean close to zero, UtJ at the right of the distribution, thus contributing to the overall skewness, and JtU at the left, then contributing to the left skewness.

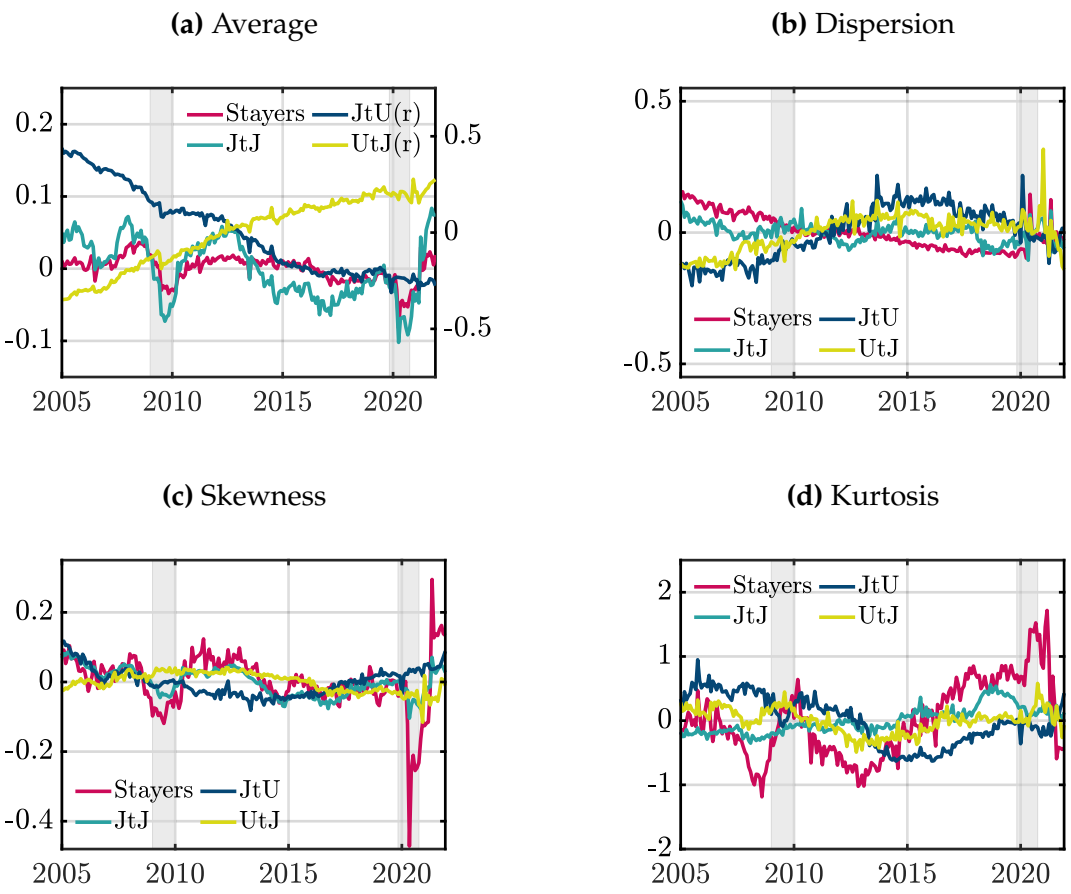
Figure 7: Distribution of twelve-month income changes conditional on labor market flow.



Note: The figure the histogram of twelve-month changes for different labor market flows: stayers, Job-to-Job, Job-to-Unemployment, and Unemployment-to-Job transitions. The figure considers data from 2015m1 to 2015m12.

While the previous description is informative, we are more interested in the evolution of these measures over the business cycle. Figure 8 shows the average, dispersion, skewness, and kurtosis evolution over time, separated by the four labor market flows. The first to note is that average growth is highly procyclical, with a pronounced fall in Chile's two recessions over that period. Notably, the dispersion of all four flows does not seem to correlate with the business cycle. On the side of higher order moments, similar to the evidence on aggregate values for the skewness and kurtosis, they are highly correlated with the business cycle, skewness for stayers, JtJ, and JtU decreases in recessions, with a very pronounced fall for stayers. We find a similar pattern for the kurtosis. This moment for stayers is the most fluctuating, increasing significantly in the recessions we show.

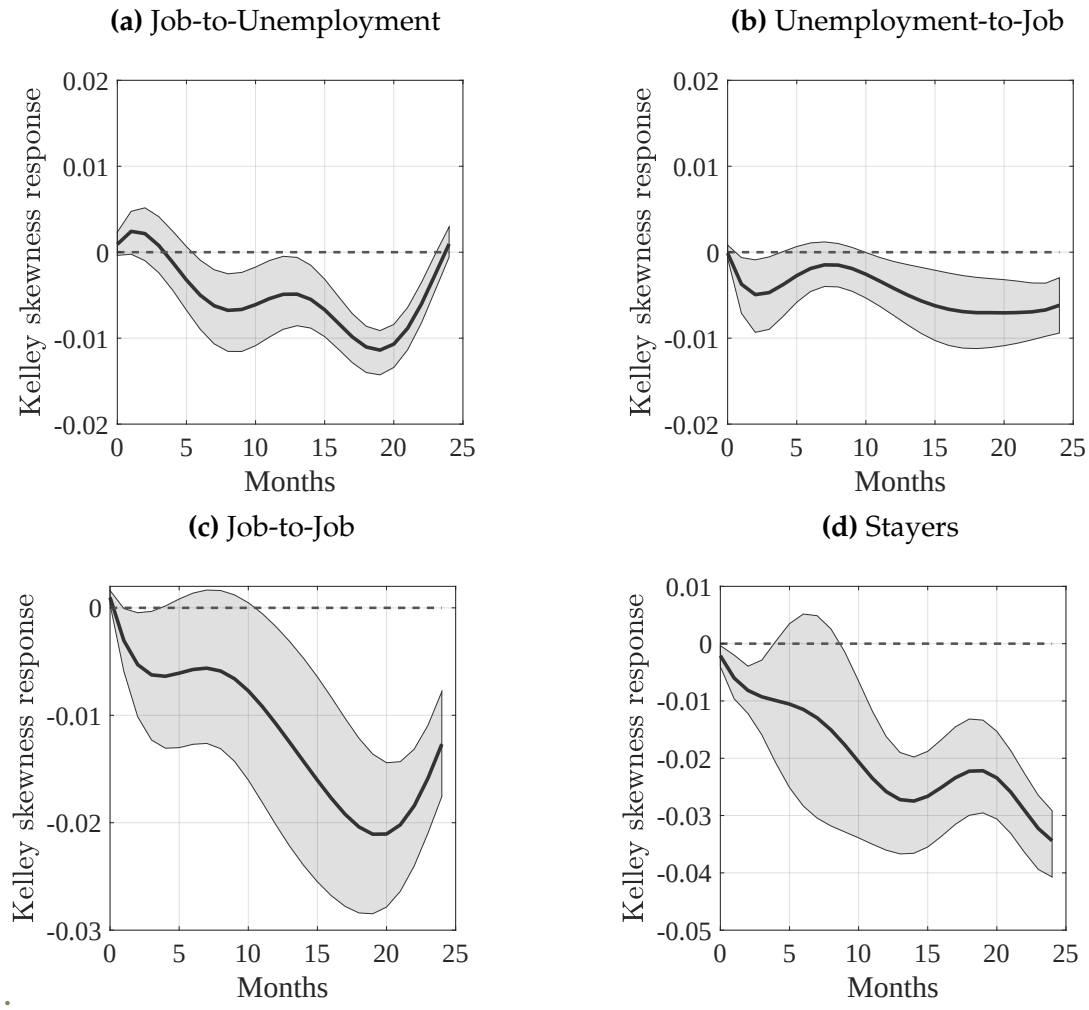
Figure 8: Time-series of moments of the twelve-month income changes conditional on labor market flows.



Note: This figure shows the time-series of the moments of the cross-sectional distribution of income growth at 12 month for different labor market flows. We show the average, dispersion, skewness, and kurtosis. Each plot shows a moment. Each figure shows the moment of different labor market flows: stayers, Job-to-Job, Job-to-Unemployment, and Unemployment-to-Job transitions. The figure considers data from 2005m9 to 2021m12. The data are detrended. Shaded areas represent recessions defined as periods with more than 6 consecutive months of y-o-y drop in GDP.

Figure 9 shows the response of the skewness of the labor market flows to a monetary policy shock. We run the same regressions in Figure 3 but with the skewness conditional on different transitions. This exercise helps us study the role of unemployment risk separated from pure income risk in response to a monetary policy shock.

Figure 9: Responses of Skewness of Labor Flows to Monetary Surprises.



Note: This figure shows the response of the Kelley skewness of the 12-months earnings growth distribution to a monetary policy shock for different labor market transitions, JtU, UtJ, JtJ and stayers.

Interestingly, all of our main transitions contribute to the aggregate negative response of the Kelley skewness. Since in response to a monetary policy shock, unemployment rises, as Figure 9a JtU risk increases. This implies a significant loss in earnings for workers, which in turn has a negative impact on skewness. Moreover, the timing is synchronized with the one of unemployment and output. The maximum fall in Kelley for JtU is lower than its aggregate, about 0.01. This result also suggests compositional effects: higher income workers become unemployed after the shock with respect normal times. Figure 9b shows the inverse transition of Figure 9a, Unemployment-to-Job. We show that transitions from unemployment to jobs contribute (but less) to the aggregate skewness. We interpret this result as the fact that unemployed workers get lower-income jobs during downturns than in booms. These two responses represent that pure unemployment risk

increases due to a monetary policy shock. This is because earnings drop due to transitions from and to unemployment, which are larger after a monetary shock.

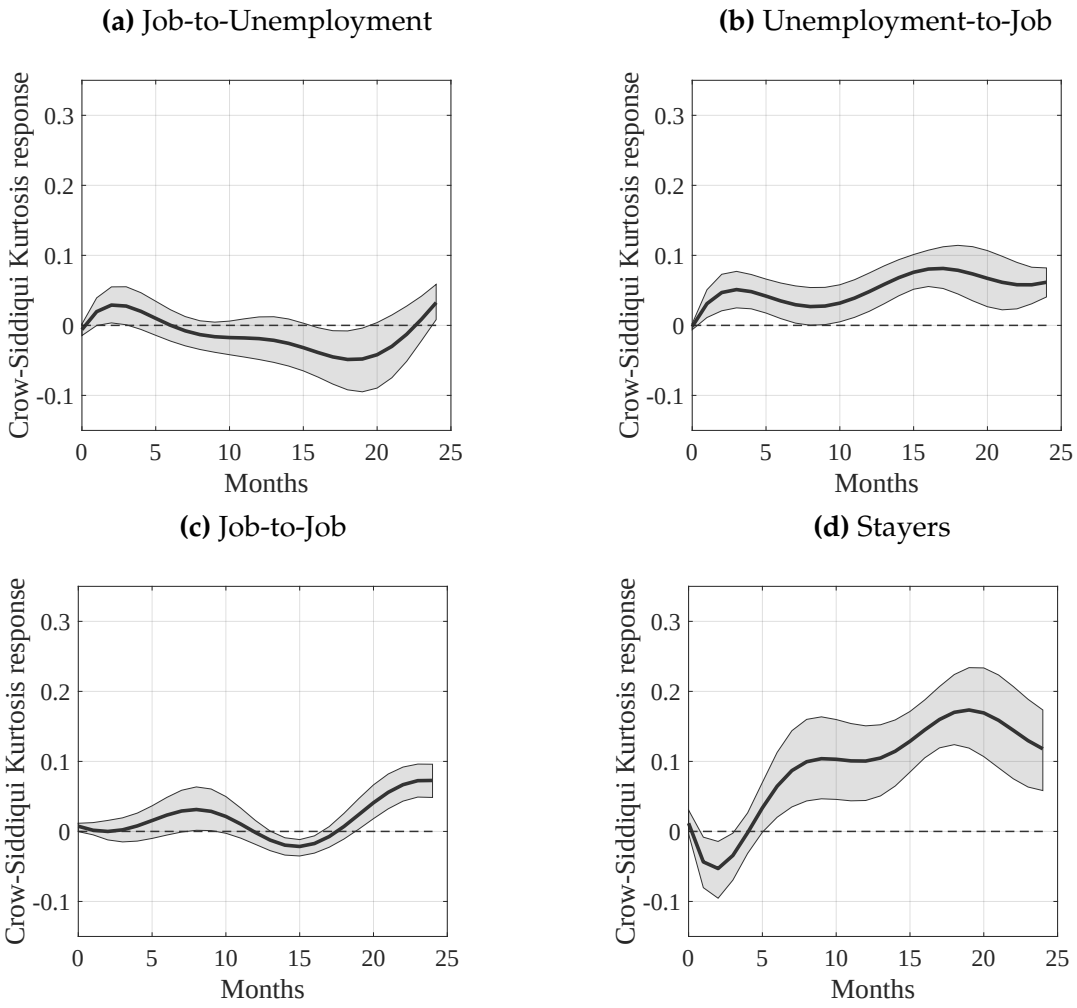
On the other hand, Figures 9c and 9d show different concepts of risk, the one coming from switching jobs and the ones from stayers in the same job. JtJ transition risk also increases in response to a monetary shock. The trough is about -0.02, so these transitions significantly contribute to the aggregate skewness, proportionally more than transitions from and to unemployment. Figure 9d surprisingly shows that the most significant contributor to aggregate skewness is the one that stayers face. We find that the trough of the response of the stayers' skewness is at -0.03, which is thrice the response of JtU risk. This means that risk is not only a phenomenon related to unemployment. People who stay in their jobs also suffer a higher risk in a downturn.

Figure 10 does the same exercise but for the kurtosis. We find that the main contributor to aggregate kurtosis is the kurtosis of stayers, followed by UtJ, with the remainder having non-significant responses. We interpret these results as evidence of lower wage rises after a monetary policy shock. Since the economy contracts, a lower mass of workers receives significant rises, thus concentrating the distribution more around its average.

For completeness, in Appendix E, we discuss the responses of the transitions to the monetary policy shock. We find that all the transitions respond as expected: JtU rates increase, UtJ rates decrease, JtJ increase, and more interestingly, the rate of stayers goes up. This last fact suggests that a monetary policy shock generates unemployment and makes the labor market more stable because workers do not switch jobs as much as without the shock.

Discussion of the results. The main empirical result of this paper, apart from showing that income risk responds significantly to monetary surprises, is that the source of income risk does not come only from unemployment, and high-income workers suffer from higher income risk in recessions. Moreover, the cyclicality of income risk is mainly due to switching and on-the-job risk. This last finding is essential–and crucial– to understanding the sources of income risk and sheds light on how we should model these kinds of problems. Thus, this result has important theoretical implications; the most direct is that workers not only face unemployment and transition risk but also face the risk when being employed, and this risk is time-varying, countercyclical, and responds to monetary surprises accordingly. This implies that income risk as idiosyncratic risk is still needed to account for aggregate fluctuations, more so for high-income workers.

Figure 10: Responses of Kurtosis of Labor Flows to Monetary Surprises.



Note: This figure shows the response of the Crow-Siddiqui kurtosis of the 12-months earnings growth distribution to a monetary policy shock for different labor market transitions, JtU, UtJ, JtJ and stayers.

5 Model

In the previous sections, we showed that the real interest rate shapes the expected labor income growth distribution, including higher-order moments. To assess the importance of this finding on the consumption-saving behavior of risk-averse households, this section presents a Heterogeneous-Agent Model (HA) with a time-varying Markov chain that captures the procyclicality of the average and the skewness of the earnings growth distribution. To capture the idea in a stylized way, we take a Huggett (1993) model and extend it to account for a time-varying income risk process in general equilibrium.

5.1 Model's Setup

Households. We assume there is a continuum of households with measure one differing in wealth $a \in \mathcal{A}$ and productivity $z \in \mathcal{Z}$. The a-dimension is the endogenous state, and the z-dimension is an exogenous time-varying idiosyncratic shock, where the idiosyncratic productivity evolves following a time-varying Markov-chain Π_t that may depend on aggregate variables. We identify each household by h, which is associated with the tuple (a, z). Households derive utility from consumption and maximizes its lifetime utility, discounting the future at a factor $0 < \beta < 1$, given by $\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u\left(c_t^h\right)$, and the instantaneous utility function takes a CRRA form $u(c_t^h) = \frac{(c_t^h)^{1-\gamma}}{1-\gamma}$, where γ is the inverse of the intertemporal elasticity of substitution. Households maximize their lifetime utility subject to the following budget and borrowing constraints:

$$\begin{aligned} c_t^h + a_{t+1}^h &= (1 + r_t)a_t^h + \tilde{Y} z_t + T_{ht}, \\ a_{ht+1} &\geq 0. \end{aligned}$$

Households save in a liquid and risk-free asset a_t^h that returns a real interest rate r_t, receive an endowment depending on her own productivity level $\tilde{Y} z_t$, where $\tilde{Y}$ is an aggregate endowment and z_t is an individual idiosyncratic shock that generates a non-degenerate distribution of income and represents the share of income each household of type z gets. T_t^h is a transfer that closes the model to satisfy equilibrium in the assets markets.

Given optimal policies $c_t^\star(z, a)$, $a'^\star_t(z, a)$ and denoting $\Psi(z, a) = Pr(z_t = z, a_{t-1} \in A)$ the probability of that combination of states at the start of date t, the distribution Ψ_t has a law of motion

$$\Psi_{t+1}(z', a') = \sum_z \Psi_{t+1}(z', a'^{\star -1}) \Pi_t(z, z') \tag{2}$$

where $a'^{\star -1}$ is the inverse of the optimal policy a. Let's define the combination of possible states as $\boldsymbol{s} = (z, a)$, so the aggregate of a variable $x_t(\boldsymbol{s})$ is given by $\int x_t(\boldsymbol{s}) \Psi(\boldsymbol{s}) d\boldsymbol{s} = X_t$. The equilibrium is defined by the market clearing conditions of the asset market $A_t(s) = B$ and the goods market given by $\tilde{Y} = C_t(s)$.

The households' optimization problem can be described by the inter-temporal problem, which is also affected by the evolution of the probability of jumping to a different productivity level in the upcoming periods.

Intertemporal Problem. To explain the mechanisms we want to highlight in this model, we write the Euler equation of each household h as follows:

$$u'(c_t^h) = \beta \sum_{z'} \Pi_t^{z'|z} u'(c_{t+1}^h)(1 + r_{t+1}) + \lambda_t^h. \tag{3}$$

Instead of using the expectation operand ($\mathbb{E}_t$), we explicitly write the Euler equation depending on the conditional transition probabilities, $\Pi_t^{z'|z}$ to explicitly account for the time-dependence of these probabilities. This feature will help us to capture, in a reduced-form way, the effect of perceived future income risk on consumption-savings behavior. For us, fluctuations in income risk are represented by an increase in the probability of getting a very large idiosyncratic shock, which implies that for a given z today, very low values of z have a higher probability of occurrence. These fluctuations in risk have a direct effect on the Euler equation. Also, in our incomplete markets model, since there is a borrowing constraint, the Euler equation does not always hold with equality and it depends on a Lagrange multiplier on the borrowing constraint condition which we denote by λ_t^h. To explore how shocks to the Markov chain affect consumption today, we can solve Equation (4) forward up to a period T to obtain

$$u'(c_0^h) = \underbrace{\beta^T \left[\prod_{s=1}^{T} \sum_{z'} \Pi_s^{z'|z}(1 + r_s) u'(c_T^{h'}) \right]}_{\text{Risk Aversion}} + \underbrace{\left[\sum_{s=1}^{T} \beta^s \Pi_s^{z'|z} \lambda_s^{h'} + \lambda_0^h \right]}_{\text{Incomplete Markets}}, \tag{4}$$

which shows that there are two channels through which income risk fluctuations affect consumption dynamics. First, the *risk aversion channel* depends on the curvature of the utility function and is present on any representative agent model. Risk aversion in our setup generates dispersion of marginal utilities along the households states; i.e., for a given state h today and possible states h' in the future, higher risk aversion disperses the values of marginal utilities. Then, with risk aversion, if the probability of a large-adverse shock increases in the future, it is more likely that consumption drops and then the marginal utility in expectations is larger, implying a cut in consumption today [9].

The second channel is the *incomplete markets channel*. This second term follows the evolution of the Lagrange multipliers of the borrowing constraint in the model. Theoretically, the closer the consumer is to the borrowing constraint, the larger the value associated with this Lagrange

[9] Prudence gives the order of magnitude of this channel.

multiplier.[10] Thus, the Lagrange multiplier is decreasing in productivity. This implies that when the future probability of having a bad draw of z increases, the incomplete market term is larger and then consumption in period zero drops.

These channels are not novel in the literature, but we want to highlight the time-varying nature of these channels. In recessions both risk aversion and incomplete market channels will play an amplification role that produces aggregate fluctuations.[11] Also, cyclicality of income risk generates strong consumption responses of high z households, since they are the most forward looking agents in these economies. This is an important fact which is absent from the heterogeneous-agent literature that only focuses on the role of high-MPC households who are not very sensitive to *expected* income fluctuations due to their financial constraints.

5.2 Calibration and Implementation

We calibrate the model to exhibit some characteristics of the Chilean economy; in particular, we followed the approach of García et al. (2024), where they match the moments of the empirical distribution with the income process in Equation (5) following the method proposed by Farmer and Toda (2017) and Tanaka and Toda (2013, 2015).

Income distribution and income risk. The estimated stochastic process is composed by two terms, a permanent and a transitory component. We estimate the parameters of the process and then discretize it to obtain a grid and a Markov chain. We assume idiosyncratic income (in logs) is given by the sum of two processes z_{1t} and z_{2t}:

$$y_t = z_{1t} + z_{2t} \tag{5}$$

where z_{it} follows

$$z_{it} = \rho_i z_{it-1} + \sigma_i \varepsilon_{it}$$

$$\varepsilon_{it} = \begin{cases} \mu_{it} \geq p_i & \sim \mathcal{N}(0,1) \\ \mu_{it} < p_i & 0 \end{cases}$$

$$\mu_{it} \sim U[0,1].$$

[10] And they are always positive.

[11] This is similar to McKay (2017b), but without assuming that the source of cyclicality of income risk arises from labor market flows.

Therefore, we estimate parameters $\{\rho_1, \rho_2, \sigma_1, \sigma_2, p_1, p_2\}$. As noted by the previous literature, the combination of these two processes returns high kurtosis (given by a $p_i \neq 0$) and can match the moments of the growth in income at lower frequencies.

Table 4 shows the estimated parameters[12]. With these parameters, we build the transition matrix to discretize them, and we consider four points for the persistent component and 12 for the transitory component.

Table 4: Parameter estimates for idiosyncratic income process.

ρ_1	ρ_2	σ_1	σ_2	p_1	p_2
0.996	0.145	0.511	0.382	0.071	0.958

The parameters are found by minimizing a loss function that takes a proposed set of parameters and computes how far we are from the desired moments.

Table 5: Empirical and estimated moments of labor earnings in Chile at a quarterly frequency.

Moment	Data	Model
Var $\log(y_t)$	0.719	0.714
Var $\Delta \log(y_t)$	0.195	0.226
Var $\Delta_{20} \log(y_t)$	0.463	0.448
Kur $\Delta \log(y_t)$	11.18	11.617
Kur $\Delta_{20} \log(y_t)$	6.21	6.076

Notes: Source: Unemployment Fund Administration, Chile.

Table 5 shows the moments of quarterly labor income for one-quarter and twenty-quarters log-change in labor income and the variance of the log of income ($\log(y_t)$). We observe that the variance increases with the lag of the difference, and these distributions have a high kurtosis, which decreases with the lag of the change. Finally, the next calibration is used in the model's parameters; $\beta = 0.987$, $\gamma = 1$, $B = 5.6$, $r_{ss} = 0.005$, and $\tilde{Y}_{ss} = 1$.

To introduce a time-varying Markov chain we consider two positive and negative skewed transition matrices that in average return the matrix associated with the process described in Equation (5). Thus, our Markov chain consists of a linear combination –defined by a weight α_t– between a High Positive Skewness (H) Markov chain and a High Negative Skewness (N), that satisfy $\Pi_{ss} = \alpha_{ss}\Pi^P_{ss} + (1 - \alpha_{ss})\Pi^N_{ss}$. In steady state we set $\alpha_{ss} = 0.5$. Out of the steady state, we

[12]The estimated permanent process with high persistence with a half-life of around 43 years (a career shock) and a low probability of occurrence: workers receive these shocks every 3.5 years. The other shock is less persistent but more likely. Households receive it almost every quarter, while its half-life is about 0.4 quarters.

assume there is a path for α_t that follows an AR(1) process and solves the transitional dynamics of the model given this shock in general equilibrium. To solve the model in general equilibrium, we find the path of interest rates that satisfies the market clearing conditions. To do so, we set a terminal time period given by T and build the system of aggregate equations $H(\boldsymbol{r})$, with $H(\cdot)$ the sequence of market clearing conditions and $\boldsymbol{r}$ the path for the interest rate. Then, we solve for the $\boldsymbol{r}^\star$ that satisfies $H(\boldsymbol{r}^\star) = 0$.

We use a solution algorithm based in the Sequence-Space Jacobian (see Auclert et al., 2021), which is to solve for the general equilibrium path of the variables in response to MIT shocks. However, our approach differs from the standard procedure because we must compute the Jacobian of the system out of the steady-state, which is more computationally demanding. For our purposes, this is necessary to capture the evolution of the transition matrix. Additionally, we have to account for the time-varying Markov chain when computing the evolution of the agents' distribution. Therefore, we solve this by a quasi-newton algorithm when solving the market clearing conditions, inspired by the Howard improvement algorithm, through which we compute the Jacobian of the system (the derivatives of $H(\cdot)$ with respect to $\boldsymbol{r}$ assuring some stability conditions.

Our exercise differs from McKay (2017a) who model income risk as the probability of having long-run unemployment. We do not want to model our problem in that way because we want to consider pure income risk, or the one suffered by people who remain in their jobs and not employment flows' risk.[13] We also depart from Bayer et al. (2019) because we want to consider the third moment of the income change distribution as well as capture all the nonlinearities of this problem since they focus on the effect of the second moment in a second-order approximation of the model.

5.3 Quantifying the importance of income risk in a non-linear HA Model

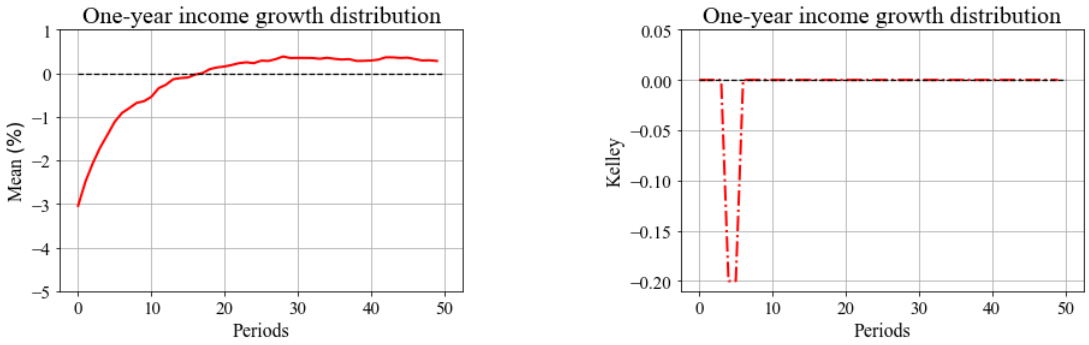
The main exercise is to study the effect of an increase in the weight of the high-risk state, and generating a decline in both the average and skewness of income growth. Thus, we assume an α_t path to create a distribution that has a year-on-year earning growth shrinkage and at the same time a decrease of the Kelley-skewness of the resulting distribution.

Figure (11) shows the mean and Kelley of the one-year income change distribution of our main exercise. Given an α_t path that moves away the Markov chain from the steady state (towards a negative skewness state), using this transition path of α_t we can simulate the quarter-on-quarter

[13] We are working on this in a companion paper.

income growth distribution, obtaining the moments' evolution. The α shock produces an on-impact effect in the average income growth distribution and a lagged response of the Kelley-skewness. In this stylized exercise, we study the effects of a 3% decline in average income growth and a decline of 0.2 in the Kelley skewness, which is an extreme case, around half the fall that was observed during the Covid-19 pandemic.

Figure 11: Simulated moments for a given α_h path to the Household's income growths.



Note: This figure shows the evolution of the simulated distribution moments (mean and Kelley skewness) of the 4-periods income growth distribution to a α_t^h path. The simulation used 350,000 points to recreate the steady state distribution $\left(\Pi(\alpha_{ss}) = \alpha_{ss}\Pi^P + (1-\alpha_{ss})\Pi^N\right)$, using the period-by-period Markov chain, given a new α_t, where $\Pi_t(\alpha_t) = \alpha_t\Pi^P + (1-\alpha_t)\Pi^N$.

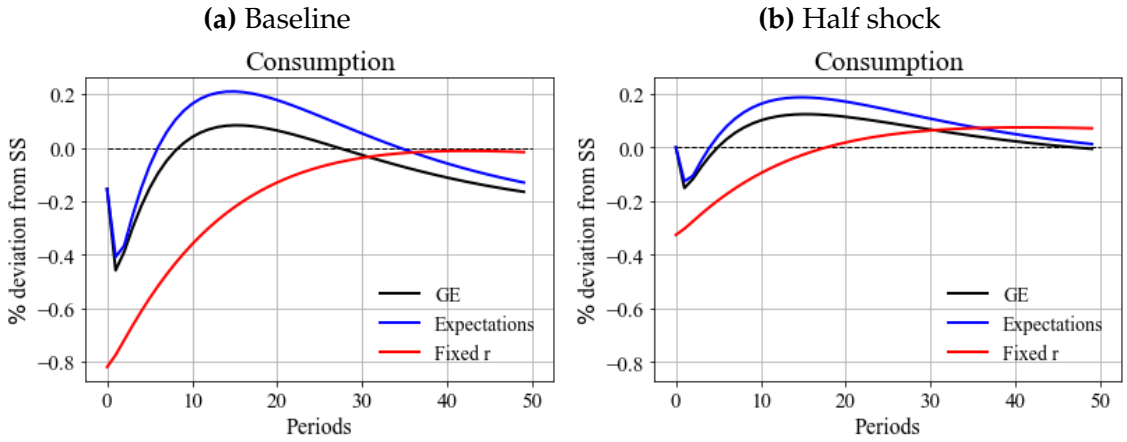
Figure 12a shows the response of consumption to this Markov chain shock. We show three different exercises. The black solid line exhibits the model's General Equilibrium (GE) response. The surge in the expected and realized income risk produces a severe negative response in aggregate consumption, with a persistent effect that lasts for several periods. One of the clear questions that comes up from this response is if the distribution of the agents is the one driving the aggregate response.

To answer that point, we use the GE prices but fixed the distribution of agents to its steady-state level. While keeping it fixed, we solve the transitional dynamic of the model and obtain the blue solid line for the aggregate consumption response. It can be seen as the expectation channel through which the income risk operates. This is so because having the distribution fixed implies that the perceived risk never materializes as it does not affect the distribution of agents. Finally, there is the partial equilibrium, it is shown in the red solid line of the figure (fixed r). It exhibits a stronger and more persistent response to the income-risk shock. This means that consumption is highly responsive to changes in expectations of future earnings.

Figure 12b shows the same exercise but with a shock that is half the one we study in Figure

12a. We can see that the responses are less than half with respect to the baseline. This implies that the responses to α_t paths are highly nonlinear, justifying the need to solve this problem with a nonlinear solver.

Figure 12: Response of Consumption to α_t^i path.



Note: The General Equilibrium path is solved using a quasi-newton algorithm, using iterative Jacobians. This is different from Auclert et al. (2021) algorithm, where they use the Jacobian matrix in the steady state.

6 Conclusion

In this paper, we show that monetary policy significantly affects labor income risk. Contractionary monetary shocks reduce average earnings growth and decrease the skewness of earnings growth. We also show that monetary policy lessens the chances of substantial earnings gains rather than increasing the probability of earnings cuts. We also find that the impact of monetary shocks on income risk varies among different worker groups, with high-income workers experiencing greater increases in risk. Additionally, workers who remain employed also show skewness declines, indicating that monetary policy's effects extend beyond transitions to and from unemployment. Our theoretical model, calibrated to our empirical results, highlights the importance of considering income risk in transmitting monetary policy. This study contributes to the literature on macroeconomic policies and worker heterogeneity, offering insights into the broader welfare costs of business cycles and the endogenous nature of income risk. Further research should highlight the role of labor market transitions and how each contributes to overall risk and and how this risk shapes the business cycle.

References

Abowd, John M, Francis Kramarz, and David N Margolis, "High wage workers and high wage firms," *Econometrica*, 1999, *67* (2), 251–333.

Aldunate, Rosario, Andrés Blanco, Andrés Fernández, Mario Giarda, and Gastón Navarro, "The Cross-Section of Labor Market Dynamics After a Foreign Shock," *mimeo*, 2025.

_ , Gabriela Contreras, and Matias Tapia, "Labor Earnings Dispersion in Chile: Decomposition, Dynamics and the Role of Firms," Working Paper 892, Central Bank of Chile November 2020.

Aruoba, Boragan, Andrés Fernández, Daniel Guzmán, Ernesto Pastén, and Felipe Saffie, "Monetary Policy Surprises in Chile: Measurement & Real Effects," Working Papers Central Bank of Chile 921, Central Bank of Chile August 2021.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 2021, *89* (5), 2375–2408.

Barnichon, Regis and Christian Brownlees, "Impulse Response Estimation by Smooth Local Projections," *The Review of Economics and Statistics*, 07 2019, *101* (3), 522–530.

Bayer, Christian, Ralph Luetticke, Lien Pham-Dao, and Volker Tjaden, "Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk," *Econometrica*, 2019, *87* (1), 255–290.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

Chan, Mons, Kory Kroft, and Ismael Mourifie, "An Empirical Framework for Matching with Imperfect Competition," 2019.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Chapter 2 Monetary policy shocks: What have we learned and to what end?," in "in," Vol. 1 of *Handbook of Macroeconomics*, Elsevier, 1999, pp. 65–148.

Debortoli, Davide and Jordi Galí, "Idiosyncratic Income Risk and Aggregate Fluctuations," *American Economic Journal: Macroeconomics*, October 2024, *16* (4), 279–310.

Farmer, Leland E. and Alexis Akira Toda, "Discretizing nonlinear, non-Gaussian Markov processes with exact conditional moments," *Quantitative Economics*, 2017, *8* (2), 651–683.

García, Benjamín, Mario Giarda, Carlos Lizama, and Ignacio Rojas, "Transmission mechanisms in HANK: An application to Chile," *Latin American Journal of Central Banking*, 2024, *5* (3), 100125.

Guiso, Luigi and Luigi Pistaferri, "The insurance role of the firm," *The Geneva Risk and Insurance Review*, 2020, *45*, 1–23.

_ , _ , and Fabiano Schivardi, "Insurance within the firm," *Journal of Political Economy*, 2005, *113* (5), 1054–1087.

Gulyas, Andreas, Matthias Meier, and Mykola Ryzhenkov, "Labor market effects of monetary policy across workers and firms," *European Economic Review*, 2024, *166*, 104756.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?," Working Paper 20913, National Bureau of Economic Research January 2015.

_ , _ , _ , and _ , "What do data on millions of US workers reveal about lifecycle earnings dynamics?," *Econometrica*, 2021, *89* (5), 2303–2339.

_ , Luigi Pistaferri, and Giovanni L Violante, "Global trends in income inequality and income dynamics: New insights from GRID," *Quantitative Economics*, 2022, *13* (4), 1321–1360.

_ , Serdar Ozkan, and Jae Song, "The nature of countercyclical income risk," *Journal of Political Economy*, 2014, *122* (3), 621–660.

Hubert, Paul and Frédérique Savignac, "The Heterogeneous Effects of Monetary Policy on Labor Income: Disentangling the Extensive and Intensive Margins," *Journal of Political Economy: Macroeconomics*, 2024, *2* (4), 687–719.

Huggett, Mark, "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of Economic Dynamics and Control*, 1993, (5-6), 953–969.

Jarociński, Marek and Peter Karadi, "Deconstructing Monetary Policy Surprises—The Role of Information Shocks," *American Economic Journal: Macroeconomics*, April 2020, *12* (2), 1–43.

Jordà, Òscar, "Estimation and inference of impulse responses by local projections," *American Economic Review*, 2005, *95* (1), 161–182.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante, "Monetary Policy According to HANK," *American Economic Review*, March 2018, *108* (3), 697–743.

McKay, Alisdair, "Time-varying idiosyncratic risk and aggregate consumption dynamics," *Journal of Monetary Economics*, 2017, *88*, 1–14.

_ , "Time-varying idiosyncratic risk and aggregate consumption dynamics," *Journal of Monetary Economics*, 2017, *88*, 1–14.

Storesletten, Kjetil, Christopher I Telmer, and Amir Yaron, "Consumption and risk sharing over the life cycle," *Journal of Monetary Economics*, 2004, *51* (3), 609–633.

Tanaka, Ken'ichiro and Alexis Akira Toda, "Discretizing Distributions with Exact Moments: Error Estimate and Convergence Analysis," *SIAM Journal on Numerical Analysis*, 2015, *53* (5), 2158–2177.

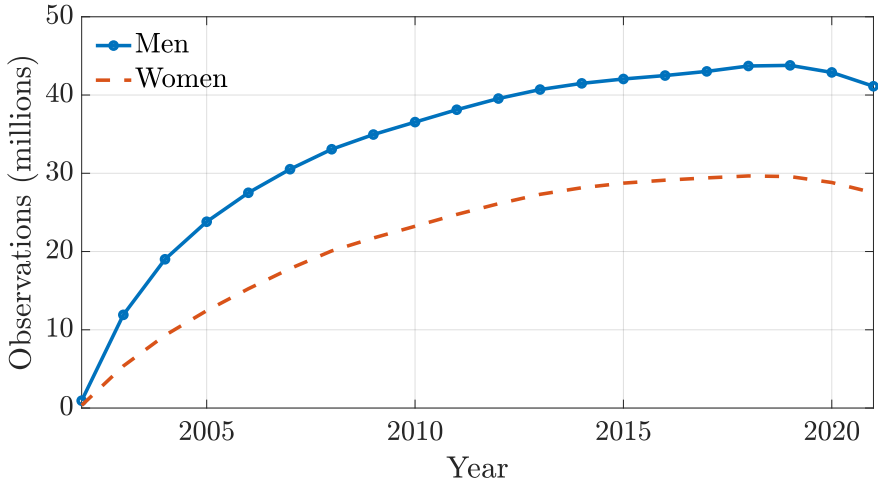
Tanaka, Ken'ichiro and Alexis Akira Toda, "Discrete approximations of continuous distributions by maximum entropy," *Economics Letters*, 2013, *118* (3), 445–450.

Online Appendix for "Monetary Shocks and Labor Income Risk with a Billion Observations"

A Unemployment Insurance Mechanism

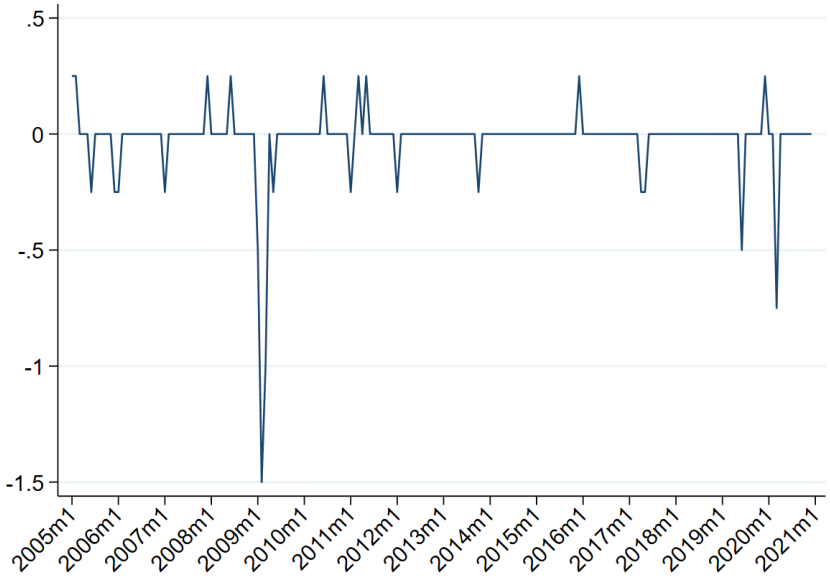
Every new employment contract signed after October 2002 is automatically enrolled in the Unemployment Insurance (UI) scheme, provided the individual is over 18 years old. For contracts signed prior to October 2002, affiliation with UI is voluntary. However, there are exceptions to who can participate in UI. Employees in the public sector, military sector, and self-employed individuals are not eligible for UI coverage. The payments from the insurance are determined based on a percentage of the employee's last ten monthly income. The percentages decrease over time: 70% for the first month, 60% for the second month, 45% for the third month, 40% for the fourth month, 35% for the fifth month, and 30% from the sixth month onwards. Financing UI is divided based on the nature of the employment contract. For long-term contracts, the employee contributes 0.6% of their monthly wage, while the employer contributes 2.4% (with 1.6% going to the employee's account and 0.8% to the solidarity fund). For short-term contracts, the employee does not contribute, but the employer contributes 3% (with 2.8% going to the employee's account and 0.2% to the solidarity fund). Any remaining funds from the UI contributions are directed towards the pension fund, and these funds are inheritable. Figure A.1 shows the evolution of yearly observations.

Figure A.1: Number of Yearly Observations



B Monetary Surprises

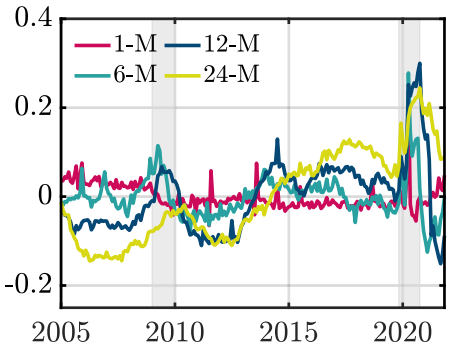
Figure B.1: Monetary Surprises



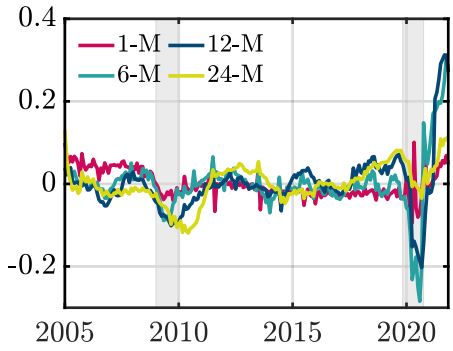
C Left and right-tail over time

Figure C.1: Demeaned Time-series of the Kelley Skewness components at different horizon income-change.

(a) P50-P10



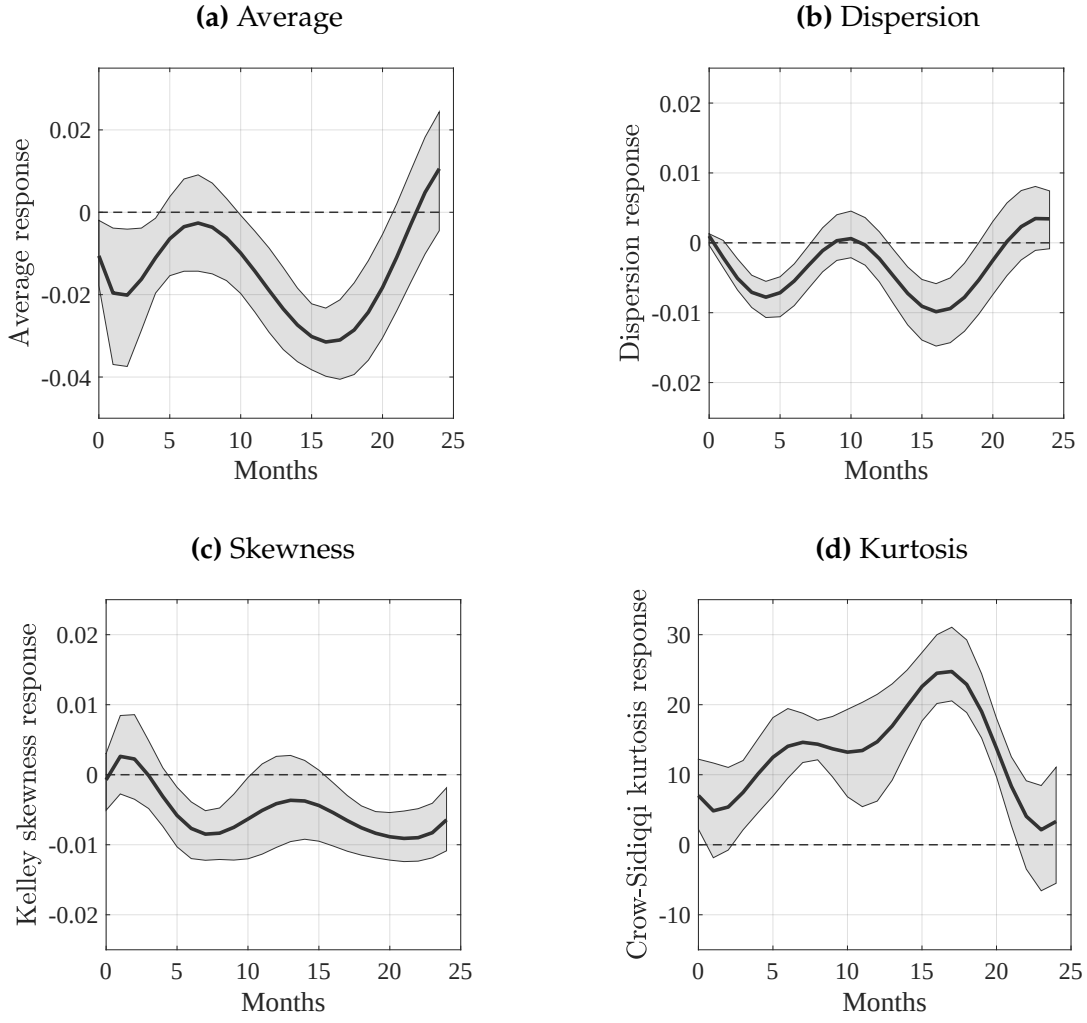
(b) P90-P50



Note: This figure shows the time-series of the P50-P10 on panel (a) and P90-P50 on panel (b) of the cross-sectional distribution of income growth at different horizons (1-, 6-, 12-, and 24-month changes distribution). The figure considers data from 2005m1 to 2021m12. The data deseasonalized (when it applies) with ARIMAX-13. Shaded areas represent recessions defined as periods with more than 6 consecutive months of y-o-y drop in GDP.

D Monthly income-change IRF

Figure D.1: Responses of moments of the income change distribution to monetary surprises.



Note: This figure shows the response of different moments of the distribution of one-month earnings growth to a monetary surprise of +25bp. The figure considers data from 2005m9 to 2019m12. Dispersion is measured as the 90th-to-10th percentiles differential. Skewness is measured as the Kelley measure, and Kurtosis is measured as the Crow-Siddiqui measure. In each Figure, we show the smoothed local projection estimates. Shaded areas are 90% confidence intervals.

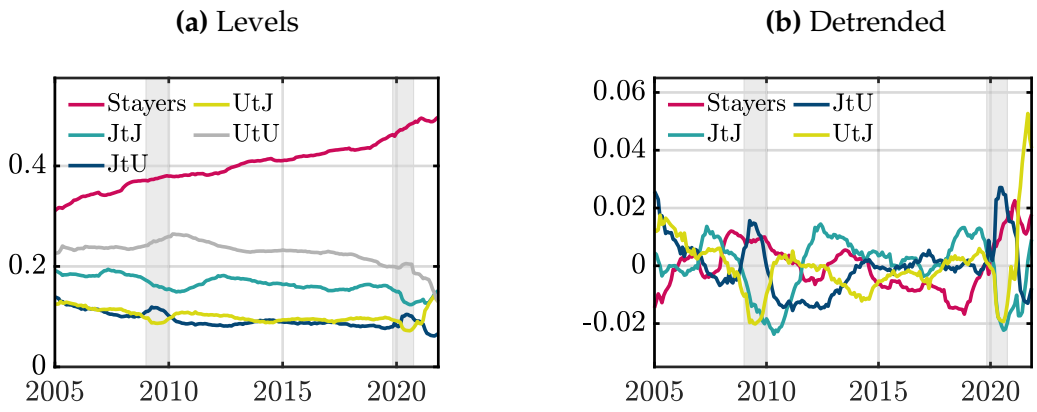
E Flows Time-series

This section analyzes the evolution of the flows over time. Even though the purpose of our paper is to understand the responses of the moments of the cross sectional distributions of income growth, it is necessary to show how the labor market flows respond to monetary policy shocks. To do so,

we build the entry (from UtJ), separation (from JtU), job-to-job, and stayers rates and analyze them over the business cycle. Separation, JtJ, and stayers rates are defined as JtU_t/L_{t-1}, JtJ_t/L_{t-1}, and S_t/L_{t-1}, with JtU_t the number of exits, JtJ_t the number of job switchers, and S_t the number of stayers in time t, and L_{t-1} the number of workers in the period before the flow, $t-1$. Whereas the entry rate is defined as UtJ_t/U_{t-1}, with UtJ_t the number of entries and U_{t-1} the number of workers with no income the period before the entry (this is a simplification of the real entry rate which is with respect to unemployment, but we do not have unemployment as itself in this data).

We show these rates in levels and the cyclical component of these flows in Figure E.1, where we include the rate of those who remained unemployed. Figure E.1b shows the cyclical component of the aforementioned ratios. From the latter, we observe that the job-to-job (JtJ) transitions have a more persistent downturn in the GFC period than the job-to-unemployment (JtU) and the unemployment-to-job (UtJ) flow, suggesting an additional transmission mechanism different than only unemployment. Additionally, the stayers' group shows an increase during the GFC period, working as a mirror of the fall of the JtJ transitions.

Figure E.1: Flows over time.

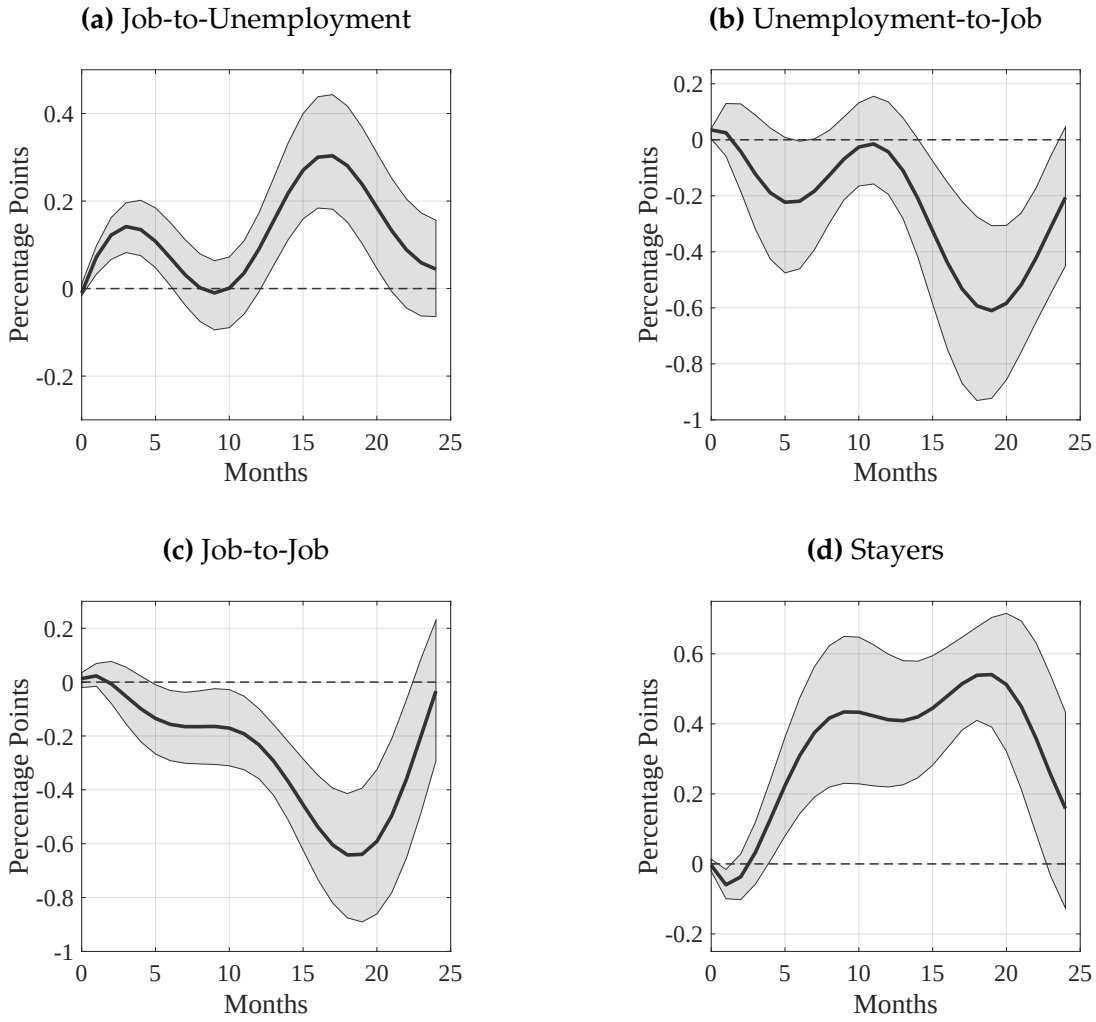


Note: This figure shows the time-series of the monements of the cross-sectional distribution of income growth at different horizons. We show the average, dispersion, skewness and kurtosis. Each plot shows a moment. Each figure shows the moment of 1-, 6-, 12-, and 24-month changes distribution. The figure considers data from 2005m1 to 2021m12. The data is detrended and deseasonalized (when it applies) with ARIMAX-13. Shaded areas represent recessions defined as periods with more than 6 consecutive months of y-o-y drop in GDP.

For completeness, in Figure E.2 we show the responses of these transition rates to monetary policy shocks. We find that all the transitions respond as expected: JtU rates increase, UtJ rates decrease, JtJ increase. More interestingly the rate of stayers goes up. This last fact suggests that a monetary policy shock not only generates unemployment but makes the labor market more stable

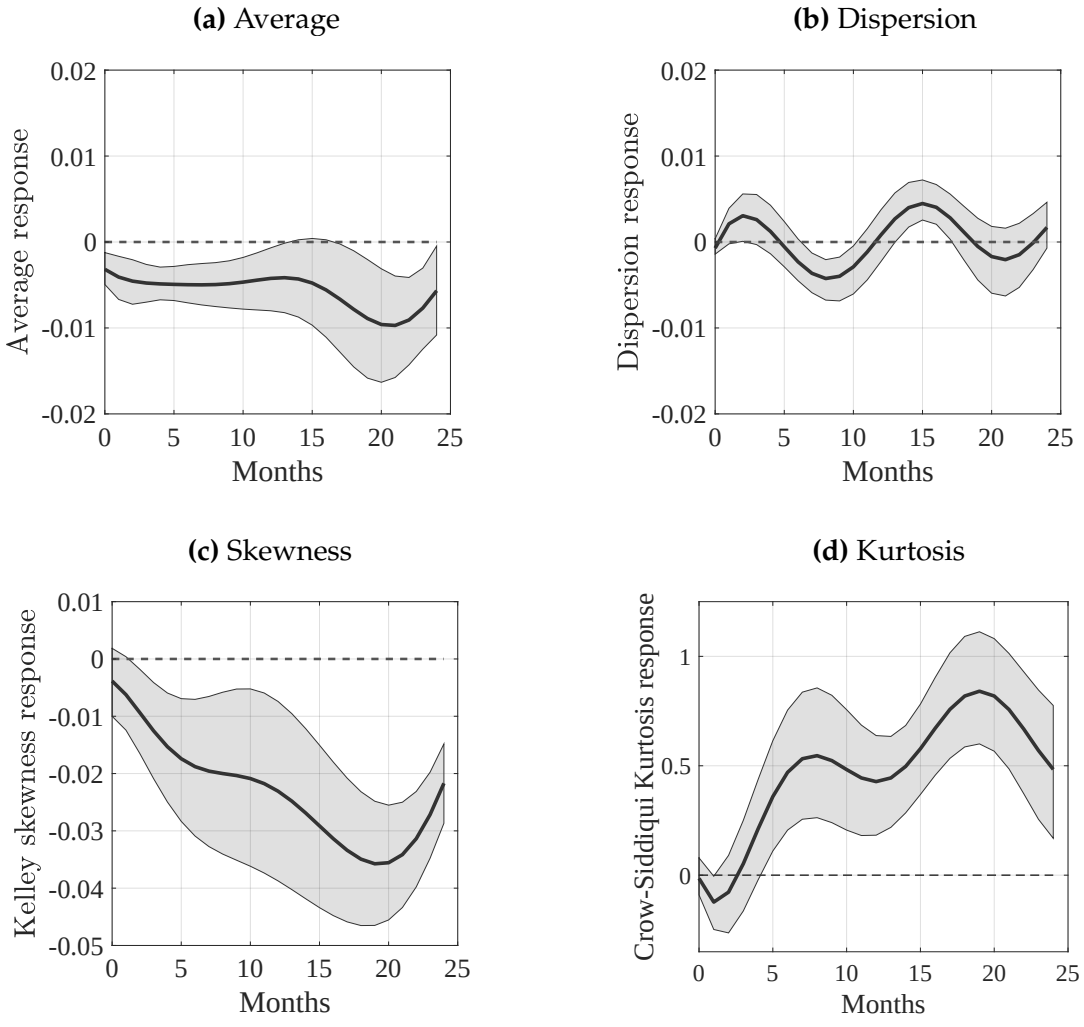
in the sense that workers do not switch jobs as much as without the shock.

Figure E.2: Responses of transition rates to monetary surprises.



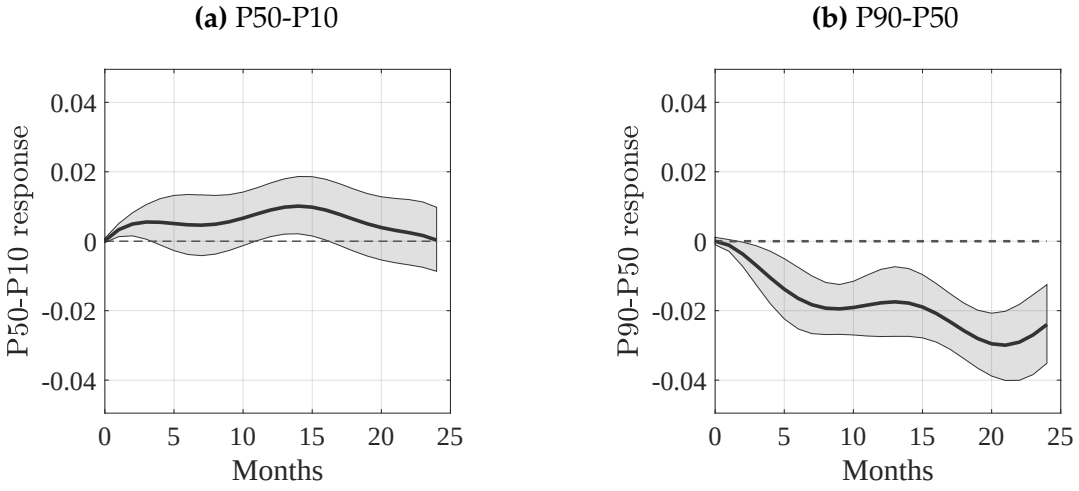
Note: This figure shows the response of the 12-months of flow rates to a monetary policy shock for different labor market transitions, JtU/L, UtJ/U, JtJ/L and stayers/L

Figure 3: Responses of moments of the income change distribution to monetary surprises.



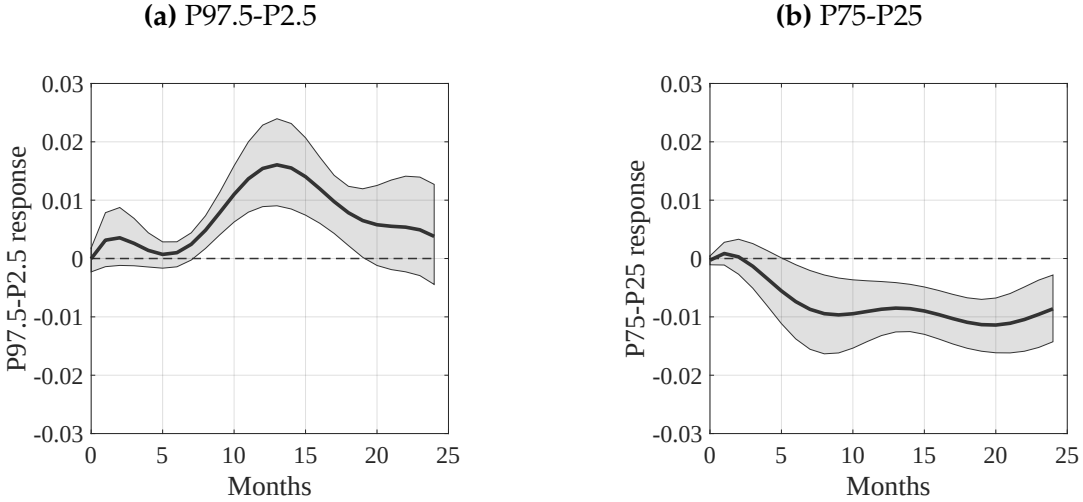
Note: This figure shows the response of different moments of the distribution of twelve-months earnings growth to a monetary surprise of +25bp. The figure considers data from 2005m9 to 2019m12. Dispersion is measured as the 90th-to-10th percentiles differential. Skewness is measured as the Kelley measure, and Kurtosis is measured as the Crow-Siddiqui measure. In each Figure, we show the smoothed local projection estimates. Shaded areas are 90% confidence intervals.

Figure 4: Response of Kelley skewness decomposition to Monetary Shocks.



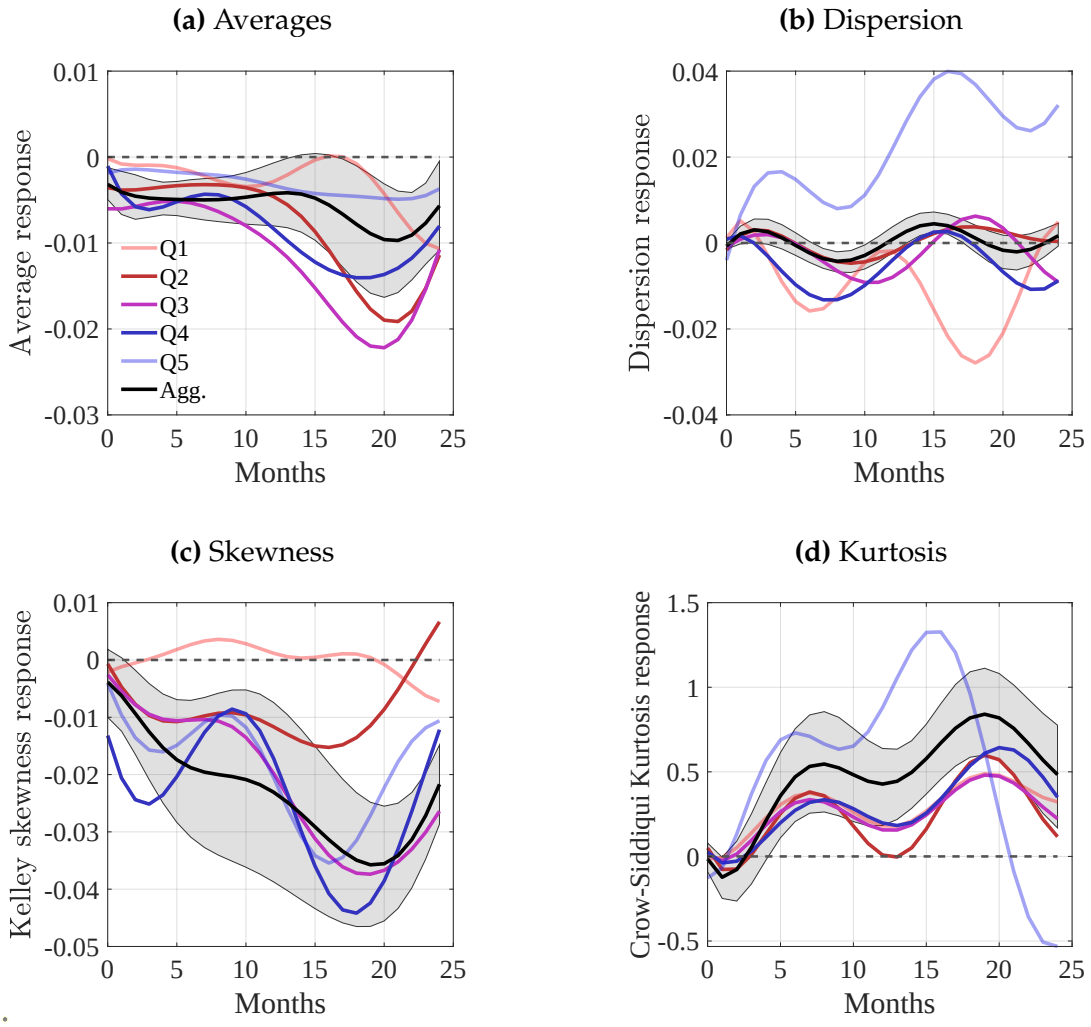
Note: This figure shows the response of the components of the Kelley skewness, $P90 - P50$ and $P50 - P10$ of the distribution of twelve-months earnings growth to a monetary surprise of +25bp. The figure considers data from 2005m9 to 2019m12. In each Figure, we show the smoothed local projection estimates. Shaded areas are 90% confidence intervals.

Figure 5: Response of Crow-Sidiqqi Kurtosis Decomposition Monetary Shocks.



Note: This figure shows the response of the components of the Cross-Siddiqui kurtosis, $P97.5 - P2.5$ and $P75 - P25$ of the distribution of twelve-months earnings growth to a monetary surprise of +25bp. The figure considers data from 2005m9 to 2019m12. In each Figure, we show the smoothed local projection estimates. Shaded areas are 90% confidence intervals.

Figure 6: Responses of Moments of the twelve-month income change distribution to monetary policy surprises by quintile of the labor income distribution.



Note: This figure shows the response of the mean, dispersion (P90-P10 differential), Kelley skewness, and Crow-Siddiqui kurtosis of the 12-months earnings growth distribution to a monetary policy shock for different groups of workers, separated by quintiles of the worker fixed effect distribution.