Fiscal Multipliers and Economic Risk[1]

From the average multiplier to the whole distribution

Francesco Frangiamore (UniPa)
Davide Furceri (IMF, CEPR)
Domenico Giannone (Johns Hopkins University, CEPR, NBER)
Giovanni Trovato (University of Rome "Tor Vergata")

Joint IMF–Bocconi University–Bank of Italy conference on
"Fiscal Policy in an Era of Uncertainty", June 4–5, 2026, Rome, Italy.

[1] The views expressed are those of the authors and should not be attributed to the IMF, its Executive Board, or IMF management.

The fiscal multiplier

The fiscal multiplier: the average effect of spending

- A central question: **how effective is fiscal policy at stimulating activity?**
- A large literature estimates the **government spending multiplier**:
  - the dollar change in output per dollar of government spending.
- Almost all of it focuses on the **average effect**.
- We build on **Ramey and Zubairy (2018)**.

Ramey and Zubairy (2018): data and specification

- Cumulative multiplier estimated directly by local projections (Jordà 2005):

$$\underbrace{\sum_{j=0}^{h} y_{t+j}}_{\text{output}} = \mu_{t+h}\Big(\underbrace{\sum_{j=0}^{h} g_{t+j}}_{\text{spending}}\Big) + \omega_{h,t+h}$$

- the **mean** μ_{t+h} — a **linear** function of future spending:

$$\sum_{j=0}^{h} y_{t+j} = \alpha_h + \beta_h \sum_{j=0}^{h} g_{t+j} + \phi_h(L) z_{t-1} + \omega_{h,t+h}$$

- y_t = GDP / potential output; g_t = government spending / potential output;
- z_{t-1}: four lags of y_t, g_t and of the shock;
- **Data:** the US, **1889–2015**, quarterly — a long history with several large wars and deep recessions.

β_h is the cumulative multiplier at horizon h — the average effect.

The data: output, spending, and the instrument

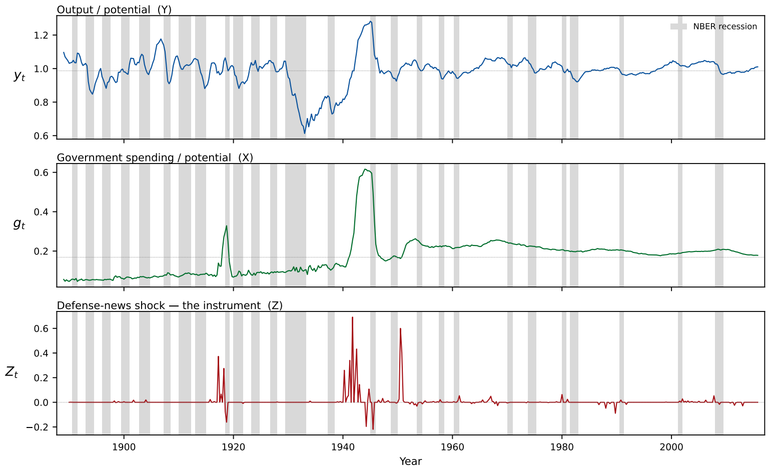

Note: US quarterly, 1889–2015. **Y**=GDP/potential (y_t); **X**=government spending/potential (g_t); **Z**=Ramey (2011) defense-news shock, the instrument (Z_t).

Shaded: NBER recessions. The instrument's variation is concentrated around major wars.

Why we cannot just regress output on spending

- Government spending is **not random** — it reacts to the economy:
  - it rises in recessions (stimulus) $\Rightarrow$ reverse causality;
  - output and spending share common drivers.

- A plain regression mixes these in, and is **biased**.

- **Instrumental variables (IV)**: an instrument Z_t that is
  (1) **relevant** — it moves spending: $Cov\left(\sum_j g_{t+j}, Z_t\right) \neq 0$;
  (2) **exogenous** — unrelated to other drivers of output: $Cov(\omega_{h,t+h}, Z_t) = 0$.

$\Rightarrow$ IV uses *only* the spending variation driven by Z_t: the **causal** multiplier.

The instrument: defense news shocks

- Instrument spending with the **Ramey (2011) defense-news shock**:
  - changes in the *expected present value of military spending*;
  - driven by **political and military events** (the build-ups to wars).

- **Exogenous** to the cycle — wars are not started to fight recessions.

- **Unanticipated** — so it shifts *expected* spending.

- **Instruments are strong** — first-stage F is large.

The classic identification strategy in the fiscal-multiplier literature.

The average effect: the Ramey–Zubairy multiplier

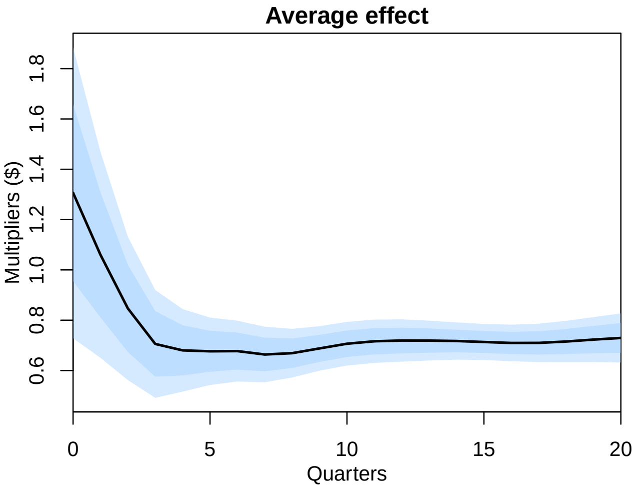

IV point estimates of the cumulative multiplier — the **average effect** of spending on output:

- above one *on impact*;
- falls quickly and settles around **0.6–0.7**.

The average multiplier is below one.

But the average is only part of the story.

Note: Cumulative IV multiplier (the average effect) over a 20-quarter horizon. Black line: point estimate; shaded areas: 68% and 90% HAC confidence intervals.

Beyond the average: economic risk

Growth-at-Risk: the predictive distribution of GDP

- Growth uncertainty is **time-varying and countercyclical:**
  - it **increases when expected (average) future growth worsens**.
- This reshapes the **predictive distribution** of GDP growth:
  - the **downside risk** varies strongly over time;
  - the **upside risk** stays roughly **stable**.
- So far, this evidence is **mostly predictive:**

$\Rightarrow$ *Our question:* does fiscal policy *affect* uncertainty — and hence the full shape of growth risk?

Our angle: a *causal* view of the distribution

- We bring the **Growth-at-Risk** lens to the **fiscal multiplier —** causally.
- Go **beyond the average** to the **whole predictive distribution** of output.

The question — a $1 cumulative rise in spending over some years:

- by how much does GDP change **on average**? (the standard multiplier)
- by how much at **each quantile** of the predictive distribution?
- in particular, what is the **uncertainty effect** — on its **variance**? (i.e. risk)

Methodology

Methodology: average *and* uncertainty effects, simultaneously

- For transparency, the **average effect** is kept exactly as in RZ:

$$\underbrace{\textstyle\sum_{j=0}^{h} y_{t+j}}_{\text{output}} = \underbrace{\hat{\mu}_{t+h}\Big(\sum_{j=0}^{h} g_{t+j}\Big)}_{\text{mean}} + \omega_{h,t+h} \qquad \hat{\mu}_{t+h}\Big(\sum_{j=0}^{h} g_{t+j}\Big) = \hat{\alpha}_h + \hat{\phi}_h(L) z_{t-1} + \hat{\beta}_h \textstyle\sum_{j=0}^{h} g_{t+j}$$

- Now **model the residual** $\omega_{h,t+h}$ as *scale* $\times$ a standardized shock:

$$\omega_{h,t+h} = \underbrace{\sigma_{t+h}\Big(\sum_{j=0}^{h} g_{t+j}\Big)}_{\text{std}} \tilde{\omega}_{h,t+h}, \qquad \tilde{\omega}_{h,t+h} \sim (\text{mean } 0, \ \text{var } 1)$$

- the **standard deviation** σ_{t+h} — a **linear** function of future spending:

$$\sigma_{t+h}\Big(\sum_{j=0}^{h} g_{t+j}\Big) = \psi_h + \eta_h(L) z_{t-1} + \gamma_h \sum_{j=0}^{h} g_{t+j}$$

$\hat{\beta}_h$: the **average effect** (unchanged from RZ). $\hat{\gamma}_h$: the **uncertainty effect** [New]

Spending can now move the mean and the std — simultaneously.

Methodology: average *and* uncertainty effects, simultaneously

- For transparency, the **average effect** is kept exactly as in RZ:

$$\underbrace{\textstyle\sum_{j=0}^{h} y_{t+j}}_{\text{output}} = \underbrace{\hat{\mu}_{t+h}\left(\sum_{j=0}^{h} g_{t+j}\right)}_{\text{mean}} + \omega_{h,t+h} \qquad \hat{\mu}_{t+h}\left(\sum_{j=0}^{h} g_{t+j}\right) = \hat{\alpha}_h + \hat{\phi}_h(L)z_{t-1} + \hat{\beta}_h \textstyle\sum_{j=0}^{h} g_{t+j}$$

- Now **model the residual** $\omega_{h,t+h}$ as *scale* $\times$ a standardized shock:

$$\omega_{h,t+h} = \underbrace{\sigma_{t+h}\left(\sum_{j=0}^{h} g_{t+j}\right)}_{\text{std}} \tilde{\omega}_{h,t+h}, \qquad \tilde{\omega}_{h,t+h} \sim (\text{mean } 0,\ \text{var } 1)$$

- the **standard deviation** σ_{t+h} — a **linear** function of future spending:

$$\sigma_{t+h}\left(\sum_{j=0}^{h} g_{t+j}\right) = \psi_h + \eta_h(L)z_{t-1} + \gamma_h \sum_{j=0}^{h} g_{t+j}$$

$\hat{\beta}_h$: the **average effect** (unchanged from RZ). $\hat{\gamma}_h$: the **uncertainty effect** [New]

Spending can now move the mean and the std — simultaneously.

Methodology: average *and* uncertainty effects, simultaneously

- For transparency, the **average effect** is kept exactly as in RZ:

$$\underbrace{\textstyle\sum_{j=0}^{h} y_{t+j}}_{\text{output}} = \underbrace{\hat{\mu}_{t+h}\Big(\sum_{j=0}^{h} g_{t+j}\Big)}_{\text{mean}} + \omega_{h,t+h} \qquad \hat{\mu}_{t+h}\Big(\sum_{j=0}^{h} g_{t+j}\Big) = \hat{\alpha}_h + \hat{\phi}_h(L) z_{t-1} + \hat{\beta}_h \textstyle\sum_{j=0}^{h} g_{t+j}$$

- Now **model the residual** $\omega_{h,t+h}$ as *scale* $\times$ a standardized shock:

$$\omega_{h,t+h} = \underbrace{\sigma_{t+h}\Big(\sum_{j=0}^{h} g_{t+j}\Big)}_{\text{std}} \, \tilde{\omega}_{h,t+h}, \qquad \tilde{\omega}_{h,t+h} \sim (\text{mean } 0,\ \text{var } 1)$$

- the **standard deviation** σ_{t+h} — a **linear** function of future spending:

$$\sigma_{t+h}\Big(\sum_{j=0}^{h} g_{t+j}\Big) = \psi_h + \eta_h(L) z_{t-1} + \gamma_h \sum_{j=0}^{h} g_{t+j}$$

$\hat{\beta}_h$: the **average effect** (unchanged from RZ). $\hat{\gamma}_h$: the **uncertainty effect** [New]

Spending can now move the mean and the std — simultaneously.

Methodology: average *and* uncertainty effects, simultaneously

- For transparency, the **average effect** is kept exactly as in RZ:

$$\underbrace{\sum_{j=0}^{h} y_{t+j}}_{\text{output}} = \underbrace{\hat{\mu}_{t+h}\left(\sum_{j=0}^{h} g_{t+j}\right)}_{\text{mean}} + \omega_{h,t+h} \qquad \hat{\mu}_{t+h}\left(\sum_{j=0}^{h} g_{t+j}\right) = \hat{\alpha}_h + \hat{\phi}_h(L) z_{t-1} + \hat{\beta}_h \sum_{j=0}^{h} g_{t+j}$$

- Now **model the residual** $\omega_{h,t+h}$ as *scale* $\times$ a standardized shock:

$$\omega_{h,t+h} = \underbrace{\sigma_{t+h}\left(\sum_{j=0}^{h} g_{t+j}\right)}_{\text{std}} \tilde{\omega}_{h,t+h}, \qquad \tilde{\omega}_{h,t+h} \sim (\text{mean } 0, \ \text{var } 1)$$

- the **standard deviation** σ_{t+h} — a **linear** function of future spending:

$$\sigma_{t+h}\left(\sum_{j=0}^{h} g_{t+j}\right) = \psi_h + \eta_h(L) z_{t-1} + \gamma_h \sum_{j=0}^{h} g_{t+j}$$

$\hat{\beta}_h$: the **average effect** (unchanged from RZ). $\hat{\gamma}_h$: the **uncertainty effect** [New]

Spending can now move the mean and the std — simultaneously.

Estimation and the quantile multiplier

Two simple steps both by IV with the defense-news shock:

1. estimate the **mean** equation $\Rightarrow \hat{\beta}_h$ (exactly the RZ multiplier);
2. regress the **absolute residuals** on the same regressors $\Rightarrow \hat{\gamma}_h$ (the scale).

The model then implies a **multiplier at every quantile** τ:

$$\boxed{\text{quantile multiplier}(\tau) \;=\; \beta_h \;+\; \gamma_h \, q_{\omega_h}(\tau)} \qquad q_{\omega_h}(\tau) \propto \Phi^{-1}(\tau)$$

- if $\gamma_h < 0$ (spending reduces uncertainty): **larger effects in the left tail** (τ small) than in the right tail.

Results

Average and uncertainty effects

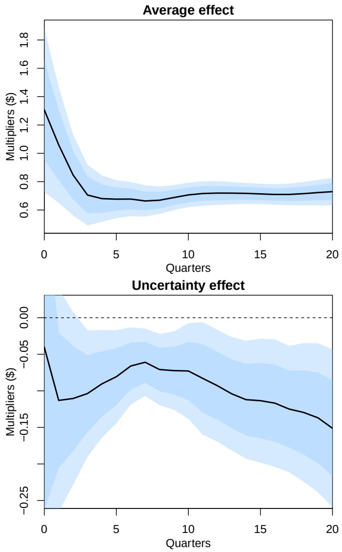

Two effects, *at the same time*:

- **Top — average effect** ↑: the multiplier (≈ 0.7), exactly Ramey–Zubairy.
- **Bottom — uncertainty effect** ↓: negative and significant.

Output *up* and uncertainty *down* — simultaneously.

Note: Black lines: point estimates; shaded areas: 68% and 90% HAC confidence bands.

Average and uncertainty effects

What it means for the predictive distribution and the quantiles

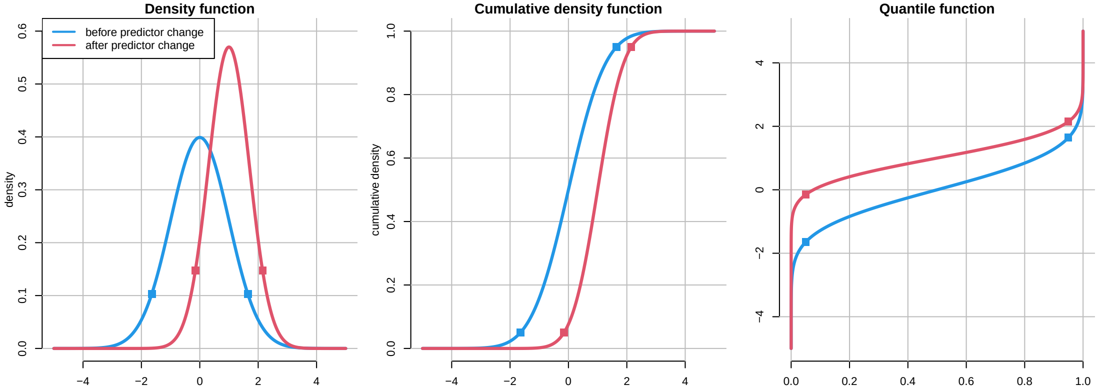

A higher *average* and lower *uncertainty* together

⇒ **left tail moves up far more than the right tail**

⇒ downside risk falls.

Average and uncertainty effects

- A fiscal expansion shifts output **up** (the usual multiplier).
- . . . and makes the distribution **narrower** — less uncertainty.
- Average *up* and uncertainty *down* reinforce in the **left tail**:
  - the **5th percentile** (bad outcome) improves a lot;
  - the **95th percentile** (good outcome) barely moves.

⇒ **Downside-risk multiplier large** (above the average); **upside multiplier small**.

⇒ Spending **mitigates** downside risk, with **limited** concerns of overheating.

The quantile multipliers make this precise.

The effects of spending across quantiles

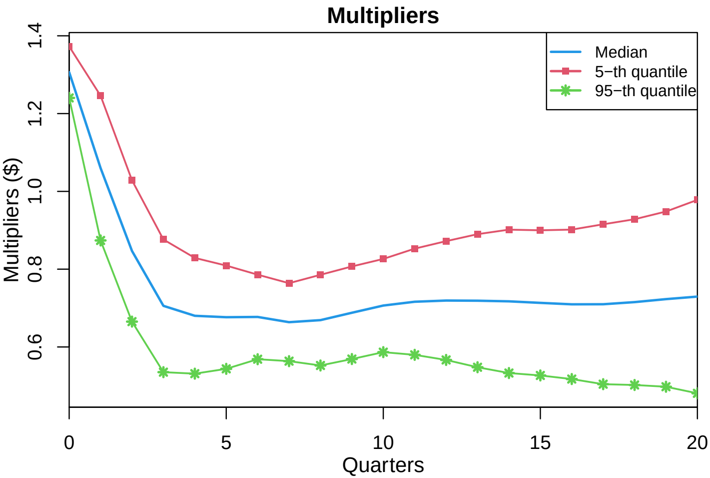

Note: Dynamic cumulative multipliers over 20 quarters at different quantiles. Blue: median; green (stars): 95th quantile; red (squares): 5th quantile. The **multiplier-at-risk (5th quantile) is about twice the 95th-quantile multiplier**, and stays above one for up to two quarters.

The effects of spending across quantiles

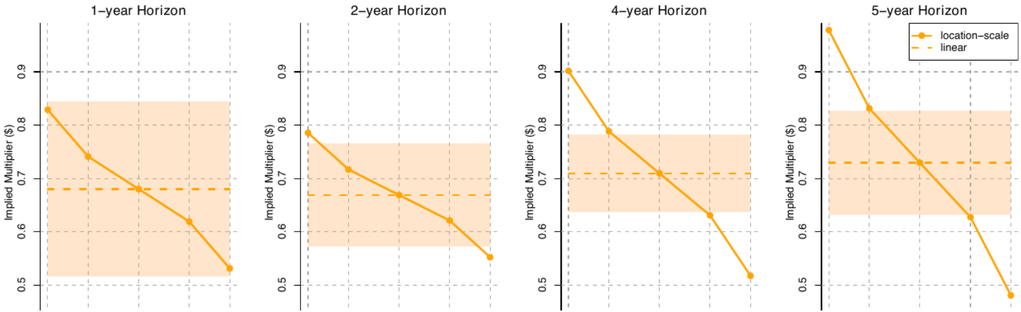

Note: Orange dashed lines: the linear RZ (2018) multiplier (the average effect) with its 90% confidence interval (shaded). Orange line with dots: multipliers at selected quantiles (x-axis) from our approach, at the 1-, 2-, 4- and 5-year horizons. The **tail multipliers lie outside the 90% band of the linear model**.

Connection with state-dependent multipliers

From quantiles to state-dependent multipliers

- Our quantile effects also speak to **state-dependent** multipliers.
- Are multipliers **larger in recessions** than in expansions?
  - Auerbach & Gorodnichenko (2012); Ramey & Zubairy (2018); ...
  - they estimate an **average effect within each state** (recession / slack).
- Key subtlety: **the state itself responds to spending**.

$\Rightarrow$ Our distributional approach captures **both** effects.

State-dependent effects: two pieces

Let μ_1, μ_0 be mean growth in recession / expansion, and p the recession probability. The state-dependent average effect splits into **two pieces**:

$$\partial_g \, \mathbb{E}[Y \mid g] \;=\; \underbrace{p\,\mu_1' + (1-p)\,\mu_0'}_{\text{(i) average effect \textit{within} the state}} \;+\; \underbrace{p'\,(\mu_1 - \mu_0)}_{\text{(ii) effect on the \textit{probability} of the state}}$$

- **(i) within-state:** the probability-weighted average of the multipliers μ_1', μ_0' — what the state-dependent literature targets.
- **(ii) probability-of-state:** spending changes the *chance* of being in each state ($p' < 0$); since $\mu_1 < \mu_0$, this piece is **positive**.

Probability distribution effects

- How does spending change the **probability of being in different states**? Estimate the probability that output falls below a threshold k, for many thresholds (Foresi & Peracchi 1995):

$$\log\left(\frac{\pi_{h,t+h,k}}{1-\pi_{h,t+h,k}}\right) = \alpha_h + \phi_h(L)z_{t-1} + \theta_h x_{h,t+h} + \varepsilon_{t+h}, \qquad \pi = \Pr(y \leq k)$$

- θ_h is the **probability distribution effect** — how spending shifts the probability of each state.
- No parametric distribution assumed — it traces the **whole CDF response** to spending.
- Estimated by **IV**: spending instrumented with the defense-news shock (as throughout).

Probability distribution effects

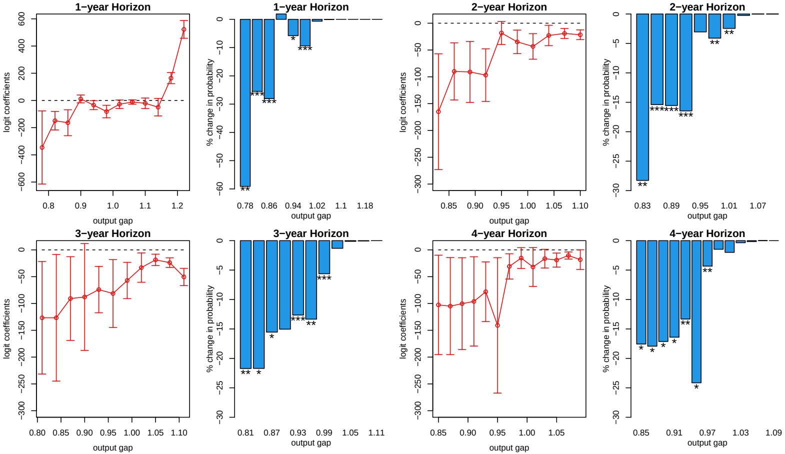

Note: Red lines with circles are the logit coefficients, while the whiskers represent the associated 90% HAC confidence intervals. The blue bars depict the corresponding marginal effects on the probability, with the stars below the bar indicating statistical significance at the 1% (***), 5% (**) and 10% (*) level.

. . . and the same two pieces explain why *uncertainty* falls

The same two pieces decompose the effect on the **variance**. Holding within-state dispersion fixed:

$$\partial_g \mathbb{V}[Y \mid g] = \underbrace{p'(1-2p)\,(\mu_1-\mu_0)^2}_{\text{regime-escape channel}} + \underbrace{2\,p(1-p)\,(\mu_1-\mu_0)(\mu_1'-\mu_0')}_{\text{mean-gap channel}}$$

- If spending **reduces the recession probability** ($p' < 0$) . . . first term < 0.
- . . . **and the multiplier is larger in recession** ($\mu_1' > \mu_0'$, so it *narrows the gap* $\mu_1 - \mu_0 < 0$) . . . second term < 0.

$\Rightarrow$ **Both channels push the variance down** — the two mechanisms are *complementary* and doubly reinforce the **reduction in uncertainty**.

This is the regime-level counterpart of our scale result ($\gamma_h < 0$).

Conclusions

Conclusions

- We go **beyond the average** to the **whole distribution** — causally.
- Spending has a **dual effect**: more output *and* less uncertainty.
- **Downside-risk multiplier sizable** — larger than the average; **upside small**.
- Fiscal policy as **insurance**: lowers *probability and severity* of bad outcomes, with limited risks of overheating.

Thanks!

francesco.frangiamore@unipa.it
dfurceri@imf.org
domenico.giannone@jhu.edu
giovanni.trovato@uniroma2.it

Annex

Annex: RZ vs. quantile during slack observations

Slack indicator	model	impact	one quarter	two quarters	three quarters	2-year integral	4-year integral
$I\{unemp_t \geq 6.5\}$ (RZ (2018) baseline)	Q5	0.23	1.30	3.76	8.66	58.62	284.05
		(1.37)	(1.25)	(1.03)	(0.88)	(0.76)	(0.90)
	Q5 + FSI	0.23	1.31	3.74	8.53	56.31	277.88
		(1.43)	(1.32)	(1.08)	(0.91)	(0.81)	(0.91)
	RZ	9.95	34.68	68.15	104.93	281.51	1104.01
		(-0.61)	(-1.92)	(-0.17)	(0.22)	(0.60)	(0.68)
$I\{unemp_t \geq 8\}$	Q5	0.16	0.85	2.38	5.46	38.09	187.74
		(1.37)	(1.25)	(1.03)	(0.88)	(0.76)	(0.90)
	Q5 + FSI	0.16	0.86	2.36	5.35	36.55	183.96
		(1.43)	(1.32)	(1.08)	(0.91)	(0.81)	(0.91)
	RZ	6.57	22.90	45.11	70.11	197.62	669.08
		(-0.44)	(-1.05)	(-0.29)	(0.16)	(0.80)	(0.76)
$I\{unemp_t \geq \tilde{unemp}_t\}$	Q5	0.29	1.70	4.97	11.37	75.01	363.00
		(1.37)	(1.25)	(1.03)	(0.88)	(0.76)	(0.90)
	Q5 + FSI	0.29	1.73	5.01	11.31	71.96	353.63
		(1.43)	(1.32)	(1.08)	(0.91)	(0.81)	(0.91)
	RZ	14.76	50.01	95.03	142.11	413.22	1887.64
		(1.30)	(0.85)	(0.61)	(0.48)	(0.52)	(0.56)
NBER recessions	Q5	0.14	0.96	2.75	6.11	44.57	273.49
		(1.37)	(1.25)	(1.03)	(0.88)	(0.76)	(0.90)
	Q5 + FSI	0.15	1.00	2.82	6.16	42.94	267.50
		(1.43)	(1.32)	(1.08)	(0.91)	(0.81)	(0.91)
	RZ	24.63	92.81	192.78	313.88	919.71	2813.84
		(0.75)	(0.82)	(0.74)	(0.66)	(0.63)	(0.67)

Note: The table shows the Sum of Squared Residuals and the multipliers in brackets, in three cases: (i) the 5-th quantile regression estimated using our approach; (ii) the 5-th quantile regression estimated using our approach, adding the Financial Stress Indicator (FSI) as control; (iii) the RZ (2018) non-linear model for the slack periods using different indicators for slack/recession indicated in the first column. $unemp_t$ is the unemployment rate, while $\tilde{unemp}_t$ indicates the Hodrick-Prescott filtered unemployment rate trend (see RZ (2018) for more details).

Annex: Alternative estimation methods

Comparison of our results with the ones obtained from the Inverse Quantile Regression (IQR) estimator by Chernozhukov and Hansen (2006) and the Smoothed Estimating Equations (SEE) estimator by Kaplan and Sun (2017), for selected horizons.

years	estimator	5-th quantile	10-th quantile	90-th quantile	95-th quantile
	FFGT	0.83	0.80	0.56	0.53
1	IQR estimator	N.C.	N.C.	N.C.	N.C.
	SEE estimator	0.82	0.80	0.48	0.21
	FFGT	0.79	0.76	0.58	0.55
2	IQR estimator	0.70	N.C.	N.C.	0.33
	SEE estimator	0.89	0.80	0.52	0.32
	FFGT	0.87	0.84	0.60	0.57
3	IQR estimator	0.88	N.C.	N.C.	0.45
	SEE estimator	0.86	0.83	0.56	0.48
	FFGT	0.90	0.86	0.56	0.52
4	IQR estimator	0.92	0.71	0.59	0.45
	SEE estimator	N.C.	0.84	0.59	0.50
	FFGT	0.98	0.92	0.54	0.48
5	IQR estimator	0.86	0.90	0.55	0.42
	SEE estimator	0.90	0.85	0.59	0.47

Note: N.C. indicates no convergence of the algorithm. FFGT = Frangiamore, Furceri, Giannone, Trovato.

Annex: spending and the probability of recession (binary)

$$\log\left(\frac{Recession_{t+h}}{1 - Recession_{t+h}}\right) = \alpha_h + \phi_h(L)z_{t-1} + \theta_h \sum_{j=0}^{h} g_{t+j} + \omega_{h,t+h}$$

	$\sum_{j=0}^{h} g_{t+j}$				
h (quarters)	0	4	8	12	16
NBER recessions	-429.23***	-6.68*	-3.40***	1.07*	0.61
	(159.37)	(3.54)	(1.05)	(0.64)	(0.52)
$unemp_t \geq 6.5$	-68.83	-6.06*	-7.27*	-1.67	-0.51
	(47.50)	(3.41)	(3.85)	(1.03)	(0.73)
$unemp_t \geq 8$	-103.63**	-7.96*	-3.24	-1.04	0.08
	(50.01)	(4.13)	(2.58)	(0.73)	(0.40)
$unemp_t \geq u\tilde{nemp}_t$	-3.98	-4.09*	-3.26*	-2.35***	0.05
	(29.92)	(2.43)	(1.93)	(0.89)	(0.43)

Note: Estimates of θ_h for different recession indicators; HAC standard errors in brackets; ***, **, * = significance at 1%, 5%, 10%. The binary-recession version of the probability piece: **higher spending significantly reduces the probability of recession** ($p' < 0$).

Annex: quantile vs. state-dependent multipliers in “slack”

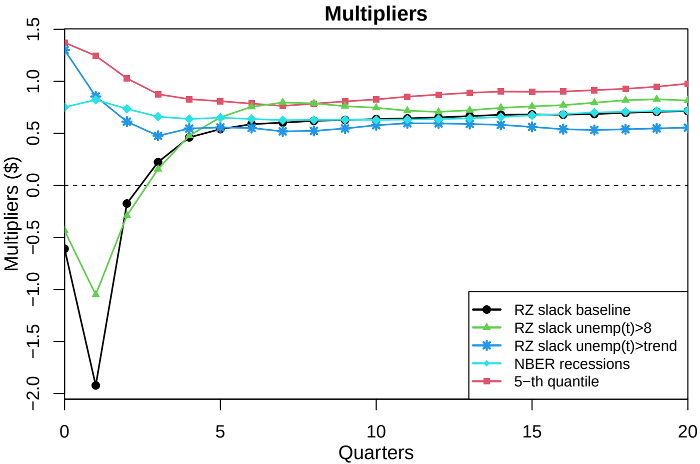

Note: RZ (2018) multipliers in slack under different slack/recession indicators (black: $unemp_t \geq 6.5$, RZ baseline; green: $unemp_t \geq 8$; blue: $unemp_t >$ HP trend; light blue: NBER). Red (squares): our 5th-quantile multiplier. **The state-dependent multiplier swings wildly with the chosen indicator; the quantile multiplier needs no ad-hoc threshold.**

Annex: Table F-statistics first stage

quarters	Location model	Scale model
0	7.33	7.33
1	8.94	8.94
2	12.86	12.86
3	16.93	16.93
4	19.22	19.22
5	19.57	19.57
6	19.26	19.26
7	19.38	19.56
8	19.24	17.77
9	18.35	17.29
10	16.86	15.43
11	15.44	14.10
12	14.23	13.83
13	13.11	13.10
14	12.06	12.35
15	11.22	11.66
16	10.73	11.09
17	10.51	10.88
18	10.47	10.84
19	10.47	10.83
20	10.43	10.88

Note: The table shows the Kleibergen-Paap rk Wald F-statistics of the first stage regressions of $\sum_{j=0}^{h} g_{t+j}$ on the Ramey (2011) fiscal news shock, computed in both the location and scale model, controlling for the same variables in the second stage regressions. The Stock and Yogo (2005) weak ID test critical values are the following: 10% maximal IV size = 16.38; 15% maximal IV size = 8.96; 20% maximal IV size = 6.66; 25% maximal IV size = 5.53.

Robustness — combined instruments

Defense-news shock combined with the Blanchard–Perotti (2002) shock.

- The Blanchard and Perotti (2002) government spending shock is identified under the assumption that government spending does not respond contemporaneously to other macroeconomic shocks within a quarter because of implementation and decision lags.
- **Exogeneity** relies on this assumption.
- There is evidence that this shock is **anticipated** (Ramey, 2011).
- **Combined instruments are stronger** – the first-stage $F \cong 70$ at short horizons (1–2 years), gradually declining to around 30 after 5 years, reflecting the greater strength of the Blanchard and Perotti (2002) shock.

Average and uncertainty effects — combined instruments

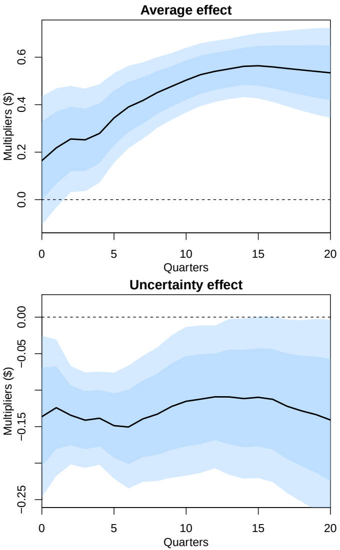

- **Top — average effect** ↑: weaker multipliers relative to the baseline, especially at shorter horizons
- **Bottom — uncertainty effect** ↓: negative and significant and stronger at short horizons relative to using the defense-news shock.

Output *up* and uncertainty *down* — simultaneously.

Note: **Top** — the *location* effect (linear/average multiplier, ≈0.6), lower than the baseline. **Bottom** — the *scale* effect (uncertainty), negative and significant. Black lines: point estimates; shaded areas: 68% and 90% HAC confidence bands.

Risk effects across quantiles — combined instruments

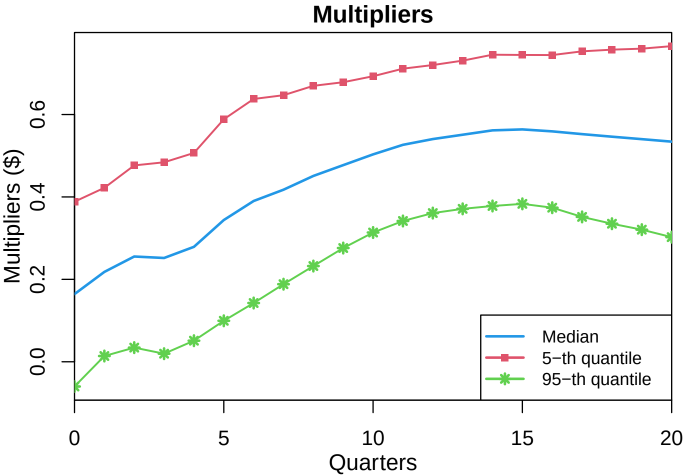

Note: Dynamic cumulative multipliers over 20 quarters at different quantiles, obtained by combining the Ramey (2011) shock with the Blanchard and Perotti (2002) shock. Blue: median; green (stars): 95th quantile; red (squares): 5th quantile. The **multiplier-at-risk (5th quantile) is about three times the 95th-quantile multiplier**, but relative to the baseline is lower.