

Change elasticity counterfactuals are not identified

Companion note to “What does it mean to change an elasticity?”

Jordan Rosenthal-Kay*
Federal Reserve Bank of San Francisco

October 2026

Brief introduction

There is growing interest in using quantitative general equilibrium models to study changes in technological elasticities: relaxing housing supply restrictions (Hsieh and Moretti, 2019) or urban costs (Rosenthal-Kay, 2026), or modeling AI as automation that raises capital’s share of income (Rockall et al., 2025; International Monetary Fund, 2025; Imas and Moll, 2026). Such “change-elasticity” counterfactuals are not generally identified by baseline data alone. Changing an elasticity while holding a coefficient fitted under a different value can move equilibrium prices and quantities in a way that depends on arbitrary units of the model calibration. More structure is needed: for example, the AI papers cited above specify task technologies. An alternative approach is to change the research question, and ask how a quantified model behaves in response to the same shocks under different elasticity regimes.

This note builds on Greaney (2026), who documents scale dependence in Hsieh and Moretti (2019) and shows why changing an elasticity requires specifying how the supply curve moves. This note also complements Rachel (2026), who shows how productivity normalization shapes the wage effects of changing capital’s share. I illustrate the housing argument in a simpler model, then show the capital-share problem through a dollars-versus-yen comparison of income paths. The housing example also relates to Azar and Giné (2026), who study worker welfare under automation when positional goods are in fixed supply.

Almost none of the content of this note is new. My goal is to formalize some tacit knowledge held by a few practitioners which bears on important, policy-relevant work now under the public eye.

A simple example: the welfare effects of income growth with a growth bottleneck

Suppose households solve

$$\max_{C,H} C^{1-\beta} H^{\beta} \quad \text{s.t.} \quad C + pH \leq M,$$

where $0 < \beta < 1$ and $M > 0$ is exogenous labor income. Housing supply is $H = Sp^{\gamma}$, where $S > 0$ is a supply shifter and $\gamma \geq 0$ is the supply elasticity. Housing is strictly inelastic when $\gamma = 0$. Supplier profits are not rebated to these households.¹

*The views in this paper do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

¹Rebating profits does not fix the units problem; see Appendix A.

Demand is $C = (1 - \beta)M$ and $pH = \beta M$. Market clearing gives the equilibrium price

$$p = \left(\frac{\beta M}{S} \right)^{1/(1+\gamma)}.$$

Indirect utility is

$$V(p, M) = B \frac{M}{p^\beta}, \quad B = (1 - \beta)^{1-\beta} \beta^\beta.$$

Holding M , S , and β fixed,

$$\ln V = \ln B + \ln M - \frac{\beta}{1+\gamma} \ln \left(\frac{\beta M}{S} \right), \quad \frac{\partial \ln V}{\partial \gamma} = \frac{\beta}{(1+\gamma)^2} \ln \left(\frac{\beta M}{S} \right) = \frac{\beta}{1+\gamma} \ln p.$$

With S fixed, raising γ lowers welfare if $p < 1$ and raises it if $p > 1$. But whether p is above or below one depends on the units.²

In practice, S is unobserved. Given β and an assumed γ , we calibrate it from observed income M_0 and price p_0 using market clearing:

$$S(\gamma) = \frac{\beta M_0}{p_0^{1+\gamma}}, \quad \frac{d \ln S(\gamma)}{d\gamma} = -\ln p_0.$$

Differentiating through this recalibration gives

$$\frac{d \ln V}{d\gamma} = \frac{\beta}{1+\gamma} \left(\ln p_0 + \frac{d \ln S(\gamma)}{d\gamma} \right) = 0.$$

Refitting S makes each supply curve pass through the observed equilibrium. I will refer to this as “anchoring” to an observed equilibrium. Changing the elasticity with this recalibration leaves the baseline equilibrium unchanged, but can change its response to shocks.

Counterfactuals in changes

Writing the model in changes, as in Dekle et al. (2008), reveals the scale dependence issue more clearly for non-marginal changes. Let $\hat{x} = x_1/x_0$ and $\Delta\gamma = \gamma_1 - \gamma_0$. Allow income, supply, and the elasticity to change, with β fixed:

$$\hat{p}\hat{H} = \hat{M}, \quad \hat{H} = \hat{S}\hat{p}^{\gamma_1} p_0^{\Delta\gamma}, \quad \hat{V} = \hat{M}\hat{p}^{-\beta}.$$

Solving for price gives

$$\hat{p} = \left(\frac{\hat{M}}{\hat{S} p_0^{\Delta\gamma}} \right)^{1/(1+\gamma_1)}.$$

The welfare change is

$$\hat{V} = \hat{M}^{1-\frac{\beta}{1+\gamma_1}} \hat{S}^{\frac{\beta}{1+\gamma_1}} p_0^{\frac{\beta\Delta\gamma}{1+\gamma_1}}.$$

With γ fixed, $\Delta\gamma = 0$, and the baseline-price term drops out. Income and supply shocks then give the same welfare change in any units. For example, doubling supply at each price, with $\gamma = 0.9$ and $\beta = 0.3$, raises welfare by $2^{0.3/1.9} - 1 = 11.57\%$.

²The units of $S = H/p^\gamma$ also change with γ . Holding S 's numerical value fixed lets its (arbitrary) units determine even the sign of the welfare effect.

Changing γ with the fitted S fixed, however, makes the answer depend on p_0 . Express the same data in yen at $\mathcal{E}$ yen per dollar, using the same exchange rate before and after the counterfactual. Income and price become $\mathcal{E}M$ and $\mathcal{E}p$, and calibration gives $S_0^{\text{yen}} = \mathcal{E}^{-\gamma_0} S_0^{\$}$. Both models fit the same data. Yet changing γ while holding each fitted S fixed gives

$$\frac{\hat{V}^{\text{yen}}}{\hat{V}^{\$}} = \mathcal{E}^{\frac{\beta \Delta \gamma}{1+\gamma_1}}.$$

Thus the same counterfactual calculation, fitted to the same baseline data, just in yen, gives a different welfare change.

For example, suppose we observe income $M_0 = \$100$ and housing price $p_0 = \$1$, with $\beta = 0.3$ and baseline elasticity $\gamma_0 = 0.9$. We calibrate the supply shifter S from these data as follows:

$$S_0 = \frac{H_0}{p_0^{\gamma_0}} = \frac{\beta M_0}{p_0^{1+\gamma_0}} = 30.$$

Now change γ to 1.8, holding income and the calibrated shifter fixed. In dollars,

$$p_1 = \left(\frac{0.3 \times 100}{30} \right)^{1/2.8} = 1 = p_0.$$

As the price is still \$1, demand is unchanged: $C_1 = C_0 = 70$ and $H_1 = H_0 = 30$. Hence $\hat{V}^{\$} = 1$. The two supply curves intersect at this price because $1^{0.9} = 1^{1.8} = 1$.

Repeat the calculation at ¥100 per dollar. The same data are $M_0 = ¥10,000$ and $p_0 = ¥100$, so calibration gives $S_0^{\text{yen}} = 0.3 \times 10,000/100^{1.9}$. Changing γ to 1.8 with this fitted S fixed gives

$$\hat{V}^{\text{yen}} = 100^{0.3(1.8-0.9)/(1+1.8)} = 1.5590.$$

Valuing counterfactual utility at baseline prices gives $v_1^{\$} = M_0^{\$} \hat{V}$. Computing in yen and converting back to dollars,

$$v_1^{\$} = \frac{¥10,000 \times 1.5590}{100} = \$155.90.$$

The welfare gain is therefore $M_0^{\$}(\hat{V} - 1) = \$55.90$, versus zero in the dollar calculation. Converting back to dollars does not fix the problem.

What to do instead

We can still shock M or S with γ fixed. We can also compare responses to the same shock across elasticities, first fitting each model to the observed equilibrium.

Consider a 10% income shock. For each $\gamma \in \{0.9, 1.8\}$, first calibrate

$$S_0(\gamma) = \frac{\beta M_0}{p_0^{1+\gamma}}.$$

Then raise income to $M_1 = 1.1M_0$, holding γ and the fitted S fixed in each model:

$$\hat{p} = 1.1^{1/(1+\gamma)}, \quad \hat{H} = 1.1^{\gamma/(1+\gamma)}, \quad \hat{V} = 1.1^{1-\beta/(1+\gamma)}.$$

With $\beta = 0.3$, welfare rises by 8.36% at $\gamma = 0.9$ and 8.88% at $\gamma = 1.8$. More elastic supply dampens the price increase. With $M_0 = \$100$, these gains are worth \$8.36 and \$8.88 at baseline prices.




Figure 1: Changing the housing supply elasticity. In (a) and (b), S stays fixed as γ rises. The price effect depends on whether baseline p is below or above one, which depends on the units. In (c), both supply curves pass through the baseline equilibrium; higher income shifts demand out. Solid blue supply has $\gamma = 0.9$ and dotted blue supply has $\gamma = 1.8$. Demand is $p = \beta M/H$.

Now express the same data in yen. Calibration gives $S_0^{\text{yen}}(\gamma) = \varepsilon^{-\gamma} S_0^{\$}(\gamma)$. The counterfactual price is

$$p_1^{\text{yen}} = \left(\frac{\beta \varepsilon M_1}{\varepsilon^{-\gamma} S_0^{\$}(\gamma)} \right)^{1/(1+\gamma)} = \varepsilon p_1^{\$}.$$

Prices multiply by ε in both periods, so $\hat{p}$, $\hat{H}$, and $\hat{V}$ are unchanged. The exchange rate also cancels when we convert the welfare gain back to dollars:

$$\frac{M_0^{\text{yen}}(\hat{V} - 1)}{\varepsilon} = M_0^{\$}(\hat{V} - 1).$$

Figure 1 shows the difference. With S fixed, raising γ raises price in (a) and lowers it in (b); welfare moves in the opposite direction. In (c), both supply curves pass through the observed equilibrium. Changing the elasticity alone does nothing, but changes the response to an income shock.

Changing capital's share of income in the neoclassical growth model

Rachel (2026) makes the corresponding point for capital-share counterfactuals: fixing a productivity normalization while changing α imposes assumptions about the gains from automation. Let's clarify the point. In the neoclassical growth model, a counterfactual must specify what happens to production on impact and then let capital adjust. For example, a TFP shock shocks production today, which has dynamic consequences. A change-elasticity counterfactual on capital's share of income needs to specify what happens to production on impact for it to make sense, and without structure, how output instantaneously adjusts to a change in α can again depend on arbitrary model units.

The model Suppose a representative household owns the capital and supplies labor inelastically, with population one and $L_t = 1$. It maximizes

$$\sum_{t=0}^{\infty} \beta^t \ln c_t, \quad c_t + k_{t+1} = y_t = A_t k_t^{\alpha}, \quad k_0 > 0.$$

Here y_t and k_t are output and capital per worker; α is capital's income share under competitive factor pricing. Capital depreciates fully each period, which lets us solve the transition exactly. The discount factor is $\beta = 0.95$, and $0 < \alpha < 1$. The Euler equation and saving rule are

$$\frac{1}{c_t} = \beta \frac{\alpha A_{t+1} k_{t+1}^{\alpha-1}}{c_{t+1}}, \quad k_{t+1} = \alpha \beta y_t, \quad c_t = (1 - \alpha \beta) y_t.$$

Fitting the baseline

If we assume the baseline is a steady state, observing y_0 is enough to fit the model given α_0 and the other parameters. In steady state, $k_0 = \alpha_0 \beta y_0$, which allows us to infer the initial capital stock k_0 from y_0 and model parameters alone. Inverting for TFP,

$$A_0 = \frac{y_0}{k_0^{\alpha_0}}.$$

We k_0 observed, we could instead fit A_0 without assuming steady state. Regardless, just like S In the housing example, A depends arbitrarily on model units, which only cancel if $\alpha_0 = 1$ (i.e., the AK model). Let's consider changing currencies from dollars to yen at $\mathcal{E}$ yen per dollar. Both y_0 and k_0 multiply by $\mathcal{E}$, so

$$A_0^{\text{yen}} = \frac{\mathcal{E} y_0^{\$}}{(\mathcal{E} k_0^{\$})^{\alpha_0}} = \mathcal{E}^{1-\alpha_0} A_0^{\$}.$$

Specifying the technology change

It's the same economy in yen or in dollars. Now, let's raise α_0 to α_1 , holding A_0 fixed. Capital k_0 cannot adjust on impact (i.e., we do not recompute it as $0.95\alpha_1 y_0$; that would fit a different baseline), but output might. Let y_{CF} denote output on impact and $\Delta\alpha = \alpha_1 - \alpha_0$. Then counterfactual output changes are

$$\frac{y_{\text{CF}}^{\$}}{y_0^{\$}} - 1 = (k_0^{\$})^{\Delta\alpha} - 1, \quad \frac{y_{\text{CF}}^{\text{yen}}}{y_0^{\text{yen}}} - 1 = (\mathcal{E} k_0^{\$})^{\Delta\alpha} - 1.$$

With $y_0 = \$0.50$ and $\alpha_0 = 0.3$, inferred capital is $\$0.1425$, or 14.25 yen at $\mathcal{E} = 100$. Raising α to 0.5 gives $0.1425^{0.2} - 1 = -32.27\%$ in dollars and $14.25^{0.2} - 1 = 70.12\%$ in yen.

Suppose instead that we refit A upon changing α , holding both baseline y_0 and k_0 fixed. Differentiating through $A(\alpha)$,

$$A(\alpha) = \frac{y_0}{k_0^{\alpha}}, \quad \frac{d \ln A(\alpha)}{d\alpha} = -\ln k_0, \quad \frac{d \ln y}{d\alpha} = \frac{d \ln A(\alpha)}{d\alpha} + \ln k_0 = 0.$$

Recalibrating ("anchoring") therefore preserves baseline production. It does not preserve marginal products or saving: with output initially unchanged, next period's capital still changes changes from $\alpha_0 \beta y_0$ to $\alpha_1 \beta y_0$.

Solving the transition

After an unexpected permanent change to A_1, α_1 at $t = 0$, the new steady state is $k^* = (\alpha_1 \beta A_1)^{1/(1-\alpha_1)}$. The entire transition is

$$k_t = k^* \left(\frac{k_0}{k^*} \right)^{\alpha_1^t}, \quad y_t = A_1 k_t^{\alpha_1}, \quad c_t = (1 - \alpha_1 \beta) y_t.$$

Changing α the wrong way. Start at a steady state with $\alpha_0 = 0.3$ and income per person $\bar{y} = \$0.50$. Then $k_0 = \alpha_0 \beta \bar{y} = \$0.1425$ and $A_0 = \bar{y}/k_0^{\alpha_0}$. Unexpectedly and permanently raise α to 0.5, holding the fitted A_0




Figure 2: Income after a change in technology. In (a), α rises from 0.3 to 0.5 with fitted A fixed. In (b), A rises permanently by 10% in two economies with different α , each starting from steady-state income of \$0.50. In (c), both technologies pass through observed $k_0 = \$0.20$, $y_0 = \$0.50$. The baseline is off steady state; we compare no change with an anchored rise in α from 0.3 to 0.5. All income is in US dollars.

fixed (again, this is wrong). Panel (a) of Figure 2 computes this experiment in dollars and in yen at $\mathcal{E} = 100$, converting the yen path back to dollars after computing it. Both calculations start from the same economy, but their income paths satisfy

$$\frac{y_t^{\text{yen}}/\mathcal{E}}{y_t^{\$}} = \mathcal{E}^{(\alpha_1 - \alpha_0) \frac{1 - \alpha_1^{t+1}}{1 - \alpha_1}}, \quad t \geq 0.$$

On impact, income is \$0.339 in the dollar calculation and \$0.851 in the yen calculation. As capital adjusts, the paths approach \$0.382 and \$2.412 per person. Income falls in the dollar calculation and rises in the yen calculation, both on impact and in the long run.

Different baseline economies, the same shock. If we recalibrate A and k under the α change, then we remain in steady state when we change α . Just like housing, now we must shock the economy to see how it behaves in different alpha regimes. Panel (b) then compares an unexpected permanent 10% increase in A across these regimes. Each starts with income of 50 cents, but the steady-state calibration assumes that they have different baseline capital stocks because their saving rates differ. To be clear: this is a comparison of productivity shocks across economies, not a change in α from the same initial capital stock (I recognize this sounds unattractive). The steady-state calibrations are

$$\bar{k}_\alpha = \alpha \beta \bar{y}, \quad \bar{A}_\alpha = \frac{\bar{y}}{\bar{k}_\alpha^\alpha}.$$

Baseline capital is \$0.1425 and \$0.2375, respectively (note that steady-state k is pinned down by $\bar{y}$ and parameters; it is not data we feed into the calibration).

Set $A_t = 1.1 \bar{A}_\alpha$ for all $t \geq 0$. Income follows

$$\frac{y_t(\alpha)}{\bar{y}} = 1.1^{\frac{1 - \alpha^{t+1}}{1 - \alpha}}.$$

Income rises by 10% on impact in both economies. Capital accumulation raises the eventual gain to 14.59% at $\alpha = 0.3$ and 21% at $\alpha = 0.5$. Computing in yen multiplies $\bar{k}_\alpha$ by $\mathcal{E}$ and $\bar{A}_\alpha$ by $\mathcal{E}^{1-\alpha}$. Every k_t, y_t, c_t then multiplies by $\mathcal{E}$. Converting back to dollars does not change anything. This is a silly counterfactual, however, because we don't want to look at the world and suppose it reflects a steady state under two different technologies (this isn't silly in the housing example, which was a static model).

The same initial capital, a change in technology. If capital k_0 were observed, recalibration would not need to assume that we are at a steady state, and therefore changing α can induce a change in the transition dynamics without a separate productivity shock (instead, think of it as a shock to the capital accumulation path).

Keeping the initial capital stock fixed and changing α , we only need to adjust A to preserve production at the baseline point (a more attractive assumption).

In panel (c), I suppose we observe y_0 and k_0 , so we don't need to assume steady state:

$$A_j = \frac{y_0}{k_0^{\alpha_j}}, \quad j \in \{0, 1\}.$$

Compare leaving $\alpha = \alpha_0$ with an unexpected permanent change to α_1 at $t = 0$. Both paths start from the same capital stock. Writing $x_t^{(j)} = k_t^{(j)}/k_0$, their transitions satisfy

$$x_0^{(j)} = 1, \quad x_{t+1}^{(j)} = \alpha_j \beta \frac{y_0}{k_0} (x_t^{(j)})^{\alpha_j}, \quad \frac{y_t^{(j)}}{y_0} = \left(\alpha_j \beta \frac{y_0}{k_0} \right)^{\frac{\alpha_j(1-\alpha_j^t)}{1-\alpha_j}}.$$

Panel (c) starts both paths at income of 50 cents and capital of 20 cents. Without a change in α , the economy is in transition, converging to its steady state: with $\alpha_0 = 0.3$, initial saving is 14.25 cents, below the capital stock. Now unexpectedly and permanently raise α to 0.5, adjusting A as above. Output on impact is still 50 cents, but saving jumps to 23.75 cents. Capital now grows instead of shrinking.

The counterfactual is the comparison of these two paths. Neither has to start in steady state, and expressing the data in yen leaves y_0/k_0 unchanged, making the comparison is scale free. "Anchoring" preserves production at the observed point, but changes investment, consumption, and subsequent output. Here we can change α without a separate productivity shock: the transition dynamics and the observation of k_0 allow us to calibrate off steady state, so the additional assumption we make is that on impact, the new technology preserves current production. This gives the structure needed for a change-elasticity counterfactual to make sense in this environment.

References

- José Azar and Mireia Giné. Higher wages, lower welfare: Automation and positional goods. Working paper, September 2026. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7503358.
- Robert Dekle, Jonathan Eaton, and Samuel Kortum. Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Staff Papers*, 55(3):511–540, 2008. URL <https://www.nber.org/papers/w13846>.
- Brian Greaney. Housing constraints and spatial misallocation: Comment. *American Economic Journal: Macroeconomics*, 18(2):409–428, 2026. URL <https://doi.org/10.1257/mac.20230141>.
- Chang-Tai Hsieh and Enrico Moretti. Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, 11(2):1–39, 2019. URL <https://pubs.aeaweb.org/doi/abs/10.1257/mac.20170388>.
- Alex Imas and Benjamin Moll. A standard task-based solow model delivers ai-driven double-digit growth. Technical supplement, September 2026. URL https://benjaminmoll.com/task_based_solow/.
- International Monetary Fund. Boosting growth in uruguay through structural reforms and artificial intelligence. Selected Issues Paper 2025/142, International Monetary Fund, 2025. URL <https://www.imf.org/-/media/files/publications/selected-issues-papers/2025/english/sipea2025142.pdf>.
- Lukasz Rachel. Does a higher capital share imply higher wages? Working paper, September 28, September 2026. URL <https://sites.google.com/site/lukasZRachel/home>.
- Emma Rockall, Marina M. Tavares, and Carlo Pizzinelli. Ai adoption and inequality. Working Paper 25/68, International Monetary Fund, 2025. URL <https://www.imf.org/-/media/files/publications/wp/2025/english/wpica2025068-print-pdf.pdf>.
- Jordan Rosenthal-Kay. Urban costs around the world. Working paper, 2026. URL <https://jrosenthalkay.github.io/>.

Appendix A: Housing supply with profits rebated

Add supplier profits to the main-text budget: $C + pH \leq M + \Pi$, where $\Pi = pH/(1 + \gamma)$. For $\gamma > 0$, this is revenue less the integral of inverse supply, assuming zero fixed costs. At $\gamma = 0$, housing is a fixed endowment and all revenue is rent. Demand is now $pH = \beta I$, where total household income is

$$I = M + \Pi = q_\gamma M, \quad q_\gamma = \frac{1 + \gamma}{1 + \gamma - \beta}.$$

Equilibrium price and the calibrated supply shifter are

$$p = \left(\frac{\beta q_\gamma M}{S} \right)^{1/(1+\gamma)}, \quad S_0 = \frac{\beta q_{\gamma_0} M_0}{p_0^{1+\gamma_0}}.$$

Replace M with total income I in the main-text calculation. Market clearing becomes $\hat{p}\hat{H} = \hat{I} = \hat{q}\hat{M}$, where $\hat{q} = q_{\gamma_1}/q_{\gamma_0}$. Welfare changes by

$$\hat{V} = (\hat{q}\hat{M})^{1 - \frac{\beta}{1+\gamma_1}} \hat{S}^{\frac{\beta}{1+\gamma_1}} p_0^{\frac{\beta\Delta\gamma}{1+\gamma_1}}.$$

For income and supply shocks with γ fixed, the proportional welfare results do not change; changing γ , however, also changes income through the profit rebate. Raising γ reduces rebated profits and therefore total income when labor income M is fixed. This does not remove the dependence on baseline price: calculating in yen rather than dollars still changes $\hat{V}$ by a factor $\mathcal{E}^{\beta\Delta\gamma/(1+\gamma_1)}$.

Return to the main-text example, with $M_0 = \$100$ and $p_0 = \$1$, and recalibrate S_0 to include the rebate. Raising γ from 0.9 to 1.8 gives $q_0 = 1.1875$, $q_1 = 1.12$, and $I_0 = \$118.75$. Hence $\hat{V}^{\$} = 0.9491$ and $\hat{V}^{\text{yen}} = 1.4797$. Valuing utility at baseline prices gives $I_0^{\$}\hat{V}$: \$112.70 from the dollar calculation and \$175.71 from the yen calculation after conversion back to dollars. The corresponding welfare gains are $-\$6.05$ and $+\$56.96$. Rebating profits does not resolve the units problem.