Estimating Individual Responses When Tomorrow Matters

Stéphane Bonhomme
University of Chicago

Angela Denis
Banco de España

EEA-ESEM Congress

August 30th, 2023

The opinions and analysis do not necessarily coincide with the opinions and analysis of the Bank of Spain or the Eurosystem

Dynamic counterfactuals

- Economists are often interested in assessing the effect of a change in the economic environment on individual decisions.

- In dynamic settings, a change in the environment involves two margins: a contemporaneous change and a change in expectations.

- We propose a regression-based method to estimate individual responses while accounting for both margins.

- We provide conditions under a structural dynamic framework that allows us to interpret average partial effects as counterfactuals.

- This semi-structural approach allows us to study dynamic counterfactuals without the need for fully specifying and estimating a structural model.

Application: income, consumption, and income expectations

• Using a standard incomplete markets model as motivation, we focus on a consumption decision rule of the form

$$Consumption_{it} = \phi_i \left(Income_{it}, IncomeBeliefs_{it}, OtherFactors_{it} \right).$$

• We study the impact of a *tax*, which affects consumption through two channels: current income $Income_{it}$, and beliefs about future income $IncomeBeliefs_{it}$.

• Empirically, we make use of subjective expectations data to learn about the agents' beliefs.

Related literature

- Our approach differs from reduced-form methods that assume the decision rule is invariant to the change (e.g., Stock, 1989; also Arellano *et al.*, 2017).

- We also differ from structural approaches since we do not specify or estimate a full structural model; e.g., Marschak (1953), Ichimura and Taber (1999, 2002), Keane and Wolpin (2002a,b), Wolpin (2013).

- We connect to the literature on identification and estimation of individual beliefs; see Dominitz and Manski (1997), Wiswall and Zafar (2015), and Chen *et al.* (2020), among others.

- Empirical regressions of decisions on elicited beliefs are common (e.g., Guiso and Parigui, 1999, Dominitz and Manski, 2007, Lochner, 2007).

Average partial effects in dynamic settings

The static case

• Consider an individual outcome y_{it} that depends on some covariates x_{it} and z_{it}.

• Suppose that, for some function g_i,

$$y_{it} = g_i(x_{it}, z_{it}) + \varepsilon_{it},$$

where ε_{it} has zero mean given x_{it} and z_{it}.

• Consider an exogenous change in x_{it}, from $x_{it} = x$ to some other value $x_{it} = x^{(\delta)}$. A standard average partial effect associated with the change in x_{it} is

$$\Delta_i^{\mathsf{APE}}(\delta, x, z) = g_i(x^{(\delta)}, z) - g_i(x, z),$$

possibly averaged across individual observations.

Limitation of the static case

• However, to interpret Δ_i^{APE} as the average change in outcomes when x_{it} changes from x to $x^{(\delta)}$, one needs to assume that the function g_i remains constant.

• This invariance assumption is often implausible in applications where dynamics matter.

• Indeed, in many settings where the current value of x_{it} changes, beliefs about future x_{it}'s (which are implicitly contained in g_i) are likely to change as well.

• For example, under a tax, both current income and beliefs about future income are generally affected.

Our approach

• Our approach to alleviate this well-known issue is to include beliefs about future x_{it} values as additional determinants of y_{it}.

• Letting π_{it} denote the subjective distribution of $x_{i,t+1}$ at time t, we postulate that, for some function ϕ_i,

$$y_{it} = \phi_i(x_{it}, \pi_{it}, z_{it}) + \varepsilon_{it},$$

where ε_{it} has zero mean given x_{it}, π_{it} and z_{it}.

• We wish to document the effects of a change from $x_{it} = x$ to $x_{it} = x^{(\delta)}$, associated with a change in beliefs from $\pi_{it} = \pi$ to $\pi_{it} = \pi^{(\delta)}$.

• Such a joint change has two distinct effects on outcomes: a contemporaneous one, and a dynamic one associated with the change in beliefs.

APEs in dynamic settings

- We define the total average partial effect, or TAPE, as

$$\Delta_i^{\text{TAPE}}(\delta, x, \pi, z) = \phi_i(x^{(\delta)}, \pi^{(\delta)}, z) - \phi_i(x, \pi, z).$$

- We then decompose the TAPE as the sum of two terms: a contemporaneous APE, where beliefs are held constant, and a dynamic APE, which solely captures the change in beliefs.

- Formally, we decompose

$$\Delta_i^{\text{TAPE}}(\delta, x, \pi, z) = \underbrace{\phi_i(x^{(\delta)}, \pi, z) - \phi_i(x, \pi, z)}_{\text{Contemporaneous}} + \underbrace{\phi_i(x^{(\delta)}, \pi^{(\delta)}, z) - \phi_i(x^{(\delta)}, \pi, z)}_{\text{Dynamic}}.$$

Interpreting average partial effects

• To interpret Δ_i^{TAPE} as the average change in outcomes when x_{it} changes from x to $x^{(\delta)}$ and π_{it} changes from π to $\pi^{(\delta)}$, one needs to assume that the function ϕ_i remains invariant in the counterfactual.

• This invariance is weaker than the assumption that g_i (without beliefs) is invariant.

• However, this is still a substantive assumption. In particular, it requires that the law of motion of beliefs remains invariant.

• In the paper, we present a structural economic framework that allows us to discuss under which conditions Δ_i^{TAPE} can be interpreted as a counterfactual effect.

Structural interpretation

Example: Consumption and income

• Consider a standard incomplete markets model of consumption and saving behavior. For simplicity, we focus on an infinite-horizon environment, as in Chamberlain and Wilson (2000).

• Household utility over log consumption is $u_i(y_{it})$. Log income x_{it} and beliefs π_{it} about $x_{i,t+1}$ are jointly first-order Markov, with ρ_i the law of motion of beliefs.

• Households can self-insure using a risk-free bond with constant interest rate r_i, and assets z_{it} follow (for $w = \exp(x)$ and $c = \exp(y)$):

$$z_{i,t+1} = (1 + r_i)(z_{it} + w_{it}) - c_{it}.$$

• Consumption is then a function of assets, income, and income beliefs

$$y_{it} = \phi\left(x_{it}, \pi_{it}, z_{it}, u_i, \beta_i, r_i, \rho_i\right) = \phi_i\left(x_{it}, \pi_{it}, z_{it}\right).$$

Example: Consumption and income

• Consider a standard incomplete markets model of consumption and saving behavior. For simplicity, we focus on an infinite-horizon environment, as in Chamberlain and Wilson (2000).

• Household utility over log consumption is $u_i(y_{it})$. Log income x_{it} and beliefs π_{it} about $x_{i,t+1}$ are jointly first-order Markov, with ρ_i the law of motion of beliefs.

• Households can self-insure using a risk-free bond with constant interest rate r_i, and assets z_{it} follow (for $w = \exp(x)$ and $c = \exp(y)$):

$$z_{i,t+1} = (1 + r_i)(z_{it} + w_{it}) - c_{it}.$$

• Consumption is then a function of assets, income, and income beliefs

$$y_{it} = \phi\left(x_{it}, \pi_{it}, z_{it}, u_i, \beta_i, r_i, \rho_i\right) = \phi_i\left(x_{it}, \pi_{it}, z_{it}\right).$$

Example: Consumption and income

• Consider a standard incomplete markets model of consumption and saving behavior. For simplicity, we focus on an infinite-horizon environment, as in Chamberlain and Wilson (2000).

• Household utility over log consumption is $u_i(y_{it})$. Log income x_{it} and beliefs π_{it} about $x_{i,t+1}$ are jointly first-order Markov, with ρ_i the law of motion of beliefs.

• Households can self-insure using a risk-free bond with constant interest rate r_i, and assets z_{it} follow (for $w = \exp(x)$ and $c = \exp(y)$):

$$z_{i,t+1} = (1 + r_i)(z_{it} + w_{it}) - c_{it}.$$

• Consumption is then a function of assets, income, and income beliefs

$$y_{it} = \phi\left(x_{it}, \pi_{it}, z_{it}, u_i, \beta_i, r_i, \rho_i\right) = \phi_i\left(x_{it}, \pi_{it}, z_{it}\right).$$

Example: Consumption and income

• Consider a standard incomplete markets model of consumption and saving behavior. For simplicity, we focus on an infinite-horizon environment, as in Chamberlain and Wilson (2000).

• Household utility over log consumption is $u_i(y_{it})$. Log income x_{it} and beliefs π_{it} about $x_{i,t+1}$ are jointly first-order Markov, with ρ_i the law of motion of beliefs.

• Households can self-insure using a risk-free bond with constant interest rate r_i, and assets z_{it} follow (for $w = \exp(x)$ and $c = \exp(y)$):

$$z_{i,t+1} = (1 + r_i)(z_{it} + w_{it}) - c_{it}.$$

• Consumption is then a function of assets, income, and income beliefs

$$y_{it} = \phi\left(x_{it}, \pi_{it}, z_{it}, u_i, \beta_i, r_i, \rho_i\right) = \phi_i\left(x_{it}, \pi_{it}, z_{it}\right).$$

Example (cont.): Tax counterfactual

• Suppose we wish to assess the impact on consumption of a proportional tax $T(w) = (1 - \lambda)w$ at time t.

• Consider a simple permanent-transitory income model where agents fully incorporate the effect of the tax into their beliefs.

• Then, the tax affects both the mean of log-income and the perceived conditional mean of future log-income.

$$\Delta_i^{\text{TAPE}}(\lambda, x, \pi, z) = \phi_i(x - log(\lambda), \pi(\eta - log(\lambda)), z) - \phi_i(x, \pi(\eta), z)$$

• In this simple model, ρ_i is not affected by the tax. However, in general, a structural interpretation of TAPE requires that ϕ_i (and hence, ρ_i) remains invariant in the counterfactual.

Estimating average partial effects

Econometric model

- We study identification and estimation in the model

$$y_{it} = \phi_i\,(x_{it}, \pi_{it}, z_{it}) + \varepsilon_{it},$$

where ε_{it} satisfies

$$\mathbb{E}[\varepsilon_{it}\,|\,x_{it}, \pi_{it}, z_{it}] = 0.$$

- Here we leave heterogeneity unrestricted and rely on large T for identification (see the paper for identification in short panels).

- If π_{it} were observed, the conditional mean $\phi_i\,(x, \pi, z)$ would be nonparametrically identified for all (x, π, z) in the support of $(x_{it}, \pi_{it}, z_{it})$.

- We assume π_{it} is parametrically identified, $\pi_{it} = \pi(\cdot;\theta_{it})$, and use subjective expectations data to learn about θ_{it}.

Estimation

- First, we estimate the belief parameters as

$$\widehat{\theta}_{it} = \underset{\theta}{\text{argmin}} \; d\left(m_{it}, m(\pi(\cdot;\theta))\right)$$

- Second, we specify, for P_r a family of functions

$$\phi_i(x,\theta,z;\alpha) = \sum_{r=1}^{R} \alpha_{ir} P_r(x,\theta,z),$$

and we estimate α_{ir} using penalized least squares regression

$$\widehat{\alpha} = \underset{\alpha}{\text{argmin}} \sum_{i=1}^{n} \sum_{t=1}^{T} \left(y_{it} - \sum_{r=1}^{R} \alpha_{ir} P_r \left(x_{it}, \widehat{\theta}_{it}, z_{it} \right) \right)^2 + \text{Pen}(\alpha),$$

using, e.g., OLS or the Lasso.

- Third, we estimate APEs by plugging in $\widehat{\theta}_{it}$ and $\widehat{\alpha}$ in the APE formulas (and use the double Lasso when plugging in Lasso estimates).

Application:
Consumption and expected income

Data

• We use the 1989, 1991, 1995 and 1998 waves of the Italian Survey of Household Income and Wealth (SHIW).

• We make use of expectations questions about income in the following year

• Our final cross-sectional sample has 7,796 household-year observations, and our panel sample has 1,646 household-year observations.

• We assume beliefs about next year log income are normally distributed.

• Three-step procedure: i) we estimate means and variances of beliefs, ii) estimate the consumption function by OLS and the Lasso, and iii) compute APEs corresponding to three tax counterfactuals.

Estimates of the consumption function (OLS)

	(1)	(2)	(3)	(4)	(5)
Mean expected log income		0.235	0.238	0.229	0.231
		(0.094)	(0.095)	(0.093)	(0.093)
(Mean expect. log income)·(Log family income)				0.104	0.104
				(0.061)	(0.061)
Var expected log income			-2.590		-2.613
			(1.876)		(1.941)
(Var expect. log income)·(Log family income)					-1.144
					(3.499)
Log family income	0.584	0.439	0.439	0.439	0.440
	(0.070)	(0.089)	(0.089)	(0.089)	(0.089)
Log family assets	0.010	0.018	0.018	0.019	0.018
	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
Household fixed effect	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
N observations	1,536	1,536	1,536	1,536	1,536
N households	768	768	768	768	768
R-squared	0.24	0.26	0.26	0.26	0.26
Pvalue F beliefs		0.01	0.03	0.02	0.05

Income tax counterfactuals: setup

• We use our framework and consumption function estimates to assess the effects of three income tax counterfactuals on consumption.

• We assume that the tax schedule is $T(w) = w - \lambda w^{1-\tau}$, where w is income. To define a baseline level of the tax, we rely on the estimates obtained by Holter *et al.* (2019) for Italy.

• In two *tax increase* counterfactuals, we increase the average tax by 10 percentage points, by decreasing λ: for one period only, and in all subsequent periods.

• In the *regressivity* counterfactual, we set $\tau = \tau^{\text{France}}$, while at the same time decreasing λ such that the tax change is neutral in terms of total tax revenue.

Average partial effects based on OLS estimates

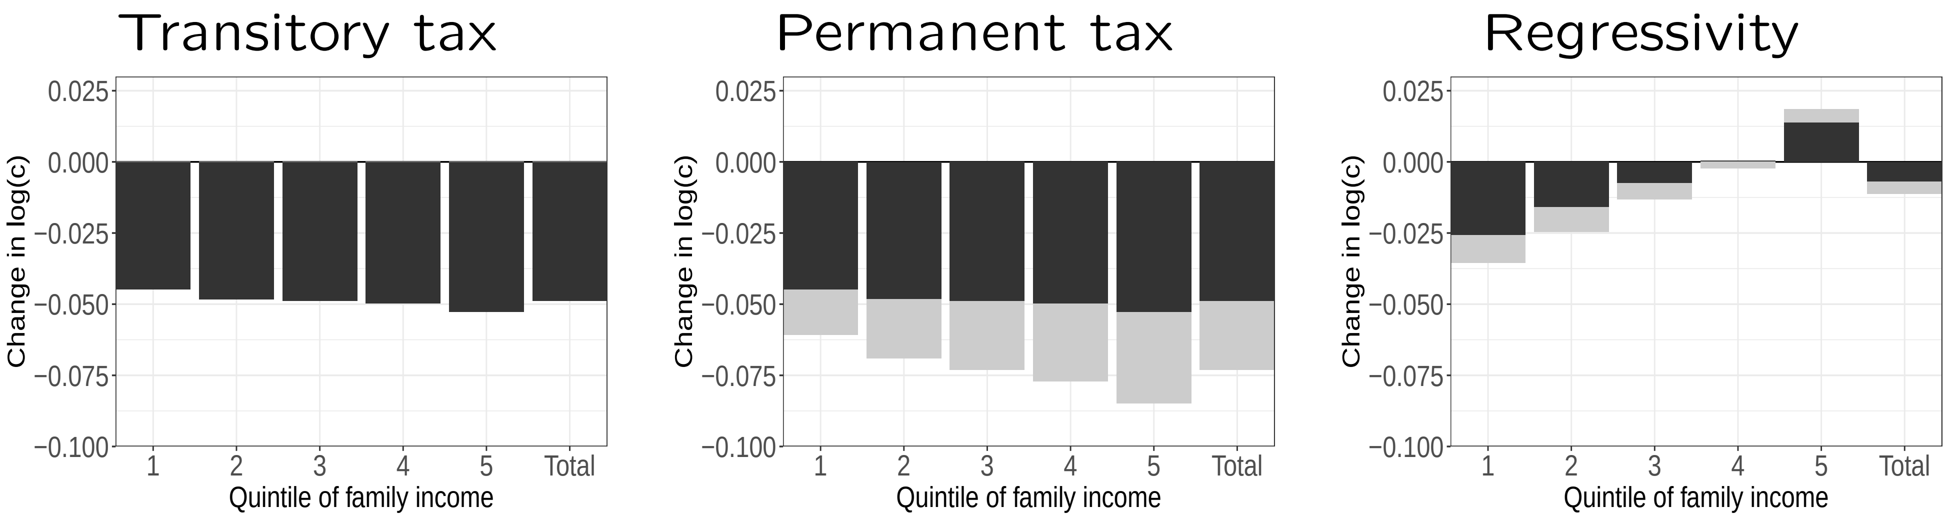

Black bars correspond to contemporaneous APE and grey bars correspond to dynamic APE. Total APE are the sums of CAPE and DAPE

Average partial effects based on the double-debiased Lasso

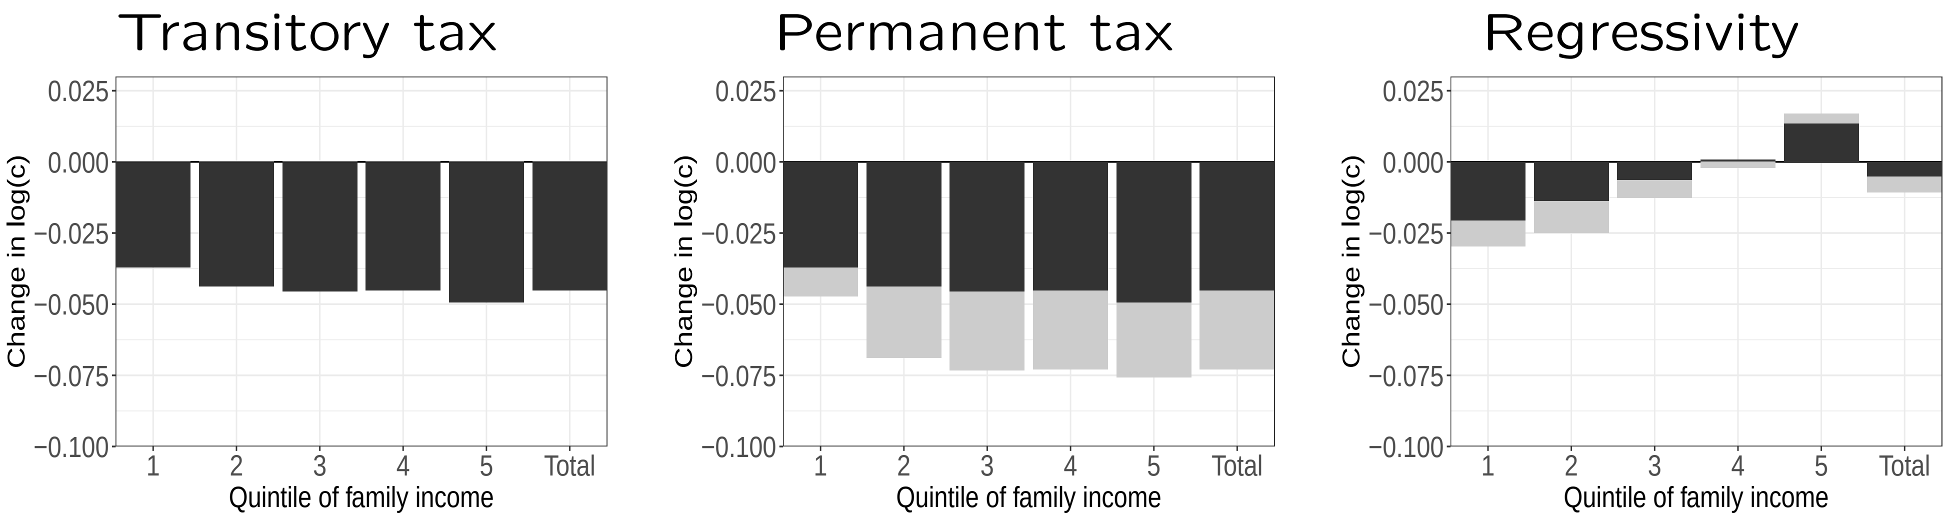

Black bars correspond to contemporaneous APE and grey bars correspond to dynamic APE. Total APE are the sums of CAPE and DAPE

Robustness

- Beliefs may be measured with error, and our ability to correct for it is limited. We perform a sensitivity analysis that suggests our estimates are not very sensitive to (a particular form of) measurement error.

- In the paper we perform other checks, using different functional forms for beliefs, studying variation over time, and varying the assumptions about assets.

Conclusion

Summary

- We propose a regression-based method to account for the role of individual expectations in assessing the impact of policies or other counterfactuals.

- We provide conditions under which APEs recover structural counterfactual effects.

- For estimation, we rely on data on elicited beliefs and propose a practical three-step method.

- Our approach does not require fully specifying and estimating a structural model, and is robust to the belief formation process, subject to the first-order Markov assumption.

Extensions

• In the absence of expectations data, our approach is still applicable provided beliefs can be estimated (e.g., under rational expectations).

• In some applications one may be interested in counterfactuals where the process of state variables changes. The framework can be applied by including state-contingent beliefs.

• In our approach, long-run beliefs ρ_i are constant in sample and invariant to the counterfactual change. This assumption can be relaxed by introducing beliefs over longer horizons.

• Lastly, extending the framework to allow beliefs to be endogenous, in the sense that past actions may shape future beliefs, will be an important task for future work.

Appendix

Compatibility with some belief formation models

- 1. Adaptive expectations: $\pi_{it} = \mathcal{N}(\mathbb{E}_{\pi_{it}}(x_{i,t+1}), \sigma_i^2)$

$$\mathbb{E}_{\pi_{it}}(x_{i,t+1}) = \mathbb{E}_{\pi_{i,t-1}}(x_{it}) + \lambda_i \left(x_{it} - \mathbb{E}_{\pi_{i,t-1}}(x_{it}) \right) + \nu_{it}.$$

- 2. Rational expectations with information Ω_{it}:

$$x_{i,t+1} = \eta_{it} + \varepsilon_{it}, \quad \Omega_{it} = \{x_i^t, \eta_i^t\},$$

where η_{it} are first-order Markov independent of ε_{it}, both normally distributed. Then, $\pi_{it} = \mathcal{N}(\eta_{it}, \sigma_i^2)$.

- 3. Rational expectations with learning: $x_{it} = \alpha_i + \varepsilon_{it}$, where agents have a normal prior on α_i, and $\varepsilon_{it} \sim \mathcal{N}(0, \sigma_{\varepsilon_i}^2)$.

Structural and semi-structural counterfactuals: simulation

	Rational expectations				Adaptive expectations			
	Structural	Semi-structural			Structural	Semi-structural		
		Linear	Quadratic	Spline		Linear	Quadratic	Spline
CAPE	-0.0163	-0.0151	-0.0150	-0.0150	-0.0122	-0.0344	-0.0191	-0.0133
DAPE	-0.0802	-0.0917	-0.0863	-0.0860	-0.0496	-0.0518	-0.0512	-0.0513
TAPE	-0.0965	-0.1068	-0.1013	-0.1010	-0.0618	-0.0863	-0.0704	-0.0646

Measurement error

• We assess measurement error in SHIW 1995–1998 when individuals are asked to distribute 100 points in a series of bins.

• Assuming individuals draw M values from $\mathcal{N}(\widehat{\mu}, \widehat{\sigma})$ and construct bootstrapped bias-corrected β-coefficients.

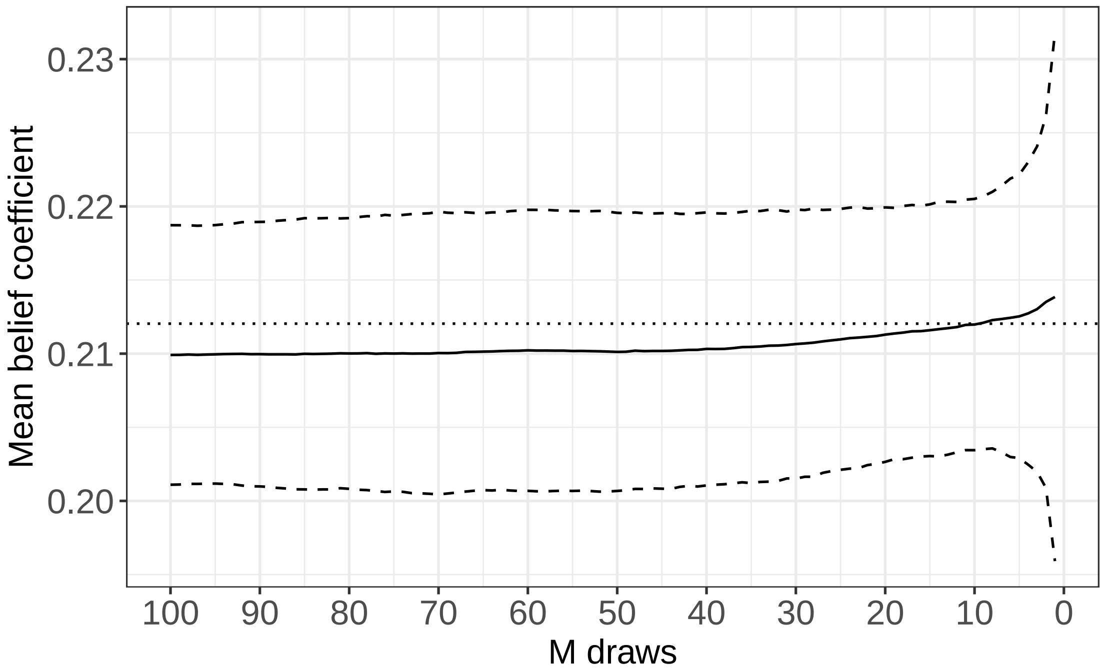

Example 2: Weather, climate and agriculture

• Consider a production function $q_{i,t+1} = g_i(x_{i,t+1}, k_{i,t+1})$, where x_{it} is the weather and k_{it} is some dynamic input.

• Suppose $k_{i,t+1} = (1 - \delta_i)k_{it} + y_{it}$, and investment has a cost $c_i(y_{it})$. The farmer decides on y_{it} after observing today's weather x_{it} and the distribution π_{it} of tomorrow's weather, but before observing $x_{i,t+1}$.

• Under appropriate conditions, investment and output are given by

$$y_{it} = \phi\left(x_{it}, \pi_{it}, k_{it}, g_i, c_i, \beta_i, \delta_i, \rho_i\right),$$
$$q_{i,t+1} = \widetilde{\phi}\left(x_{i,t+1}, x_{it}, \pi_{it}, k_{it}, g_i, c_i, \beta_i, \delta_i, \rho_i\right).$$

• This dynamic model allows one to study farmers' adaptation to a change in the weather process (e.g., Dell *et al.*, 2014, Burke and Emerick, 2016, Keane and Neal, 2020, Shrader, 2021).