Shock Propagation in Dynamic (S,s) Economies:

A State Space Approach

Miguel Bandeira[*] and Laura Castillo-Martínez[†]

June 2023

PRELIMINARY AND INCOMPLETE. PLEASE DO NOT CIRCULATE.

Abstract

This paper introduces a general statistical framework to study aggregate shock propagation in a dynamic (S,s) economy. The proposed framework enables researchers to combine micro and macro data to estimate all the parameters in this economy by maximum likelihood or Bayesian methods, estimate the cross-sectional distributions latent states over time, estimate impulse responses to aggregate shocks of any function of the cross-sectional distribution of lumpy variables given any initial distribution of state gaps and to forecast any function of the cross-sectional distribution of lumpy variables. As an application of our framework we study the propagation of monetary policy shocks in a random menu cost economy.

[*]Insper and Centre For Macroeconomics (CFM). Email: mbandeira@insper.edu.br.
[†]Duke University and CEPR. Email: l.castillo-martinez@duke.edu.

1. Introduction

Many decisions by economic agents are characterized by long periods of inaction followed by large adjustments. This type of pattern has been empirically documented using micro data for prices, plant-level capital, firm-level employment, inventory orders, individual cash withdrawals, portfolio management decisions, individual consumption of durable goods, among others. Economists have often rationalized this behavior by incorporating microeconomic adjustment frictions, in particular, fixed adjustment costs, to agent optimization problems. The resulting decision rules are known as (*S*,*s*) rules, with adjustment only taking place whenever the gap between the variable of interest and its frictionless counterpart is large enough. Despite the pervasiveness of lumpiness in the data, understanding the implications of (*S*,*s*) rules for the propagation of macroeconomic shocks is still an ongoing question.

This paper introduces a general statistical framework to study aggregate shock propagation in an economy with microeconomic adjustment frictions. Our approach imposes minimal requirements on available data and shock structure. Regarding the former, we assume the time series of aggregate variables as well as a panel of agents' choices, including the lumpy variable that follows a two-sided (*S*,*s*) rule, are observable to the econometrician. Regarding the latter, we assume that the frictionless value of the lumpy variable is subject to shocks that are conditionally iid across agents and over time. Given this setup, the proposed framework enables researchers to combine micro and macro data to (*i*) estimate all the parameters in this economy by maximum likelihood or Bayesian methods, (*ii*) estimate the cross-sectional distributions latent states over time, (*iii*) estimate impulse responses to aggregate shocks of *any* function of the cross-sectional distribution of lumpy variables given *any* initial distribution of state gaps and (*iv*) to forecast any function of the cross-sectional distribution of lumpy variables.

At the core of our framework lies a state space representation of the data generating process of the dynamic (*S*,*s*) economy. We allow for two idiosyncratic latent states, the frictionless value of the lumpy variable and an indicator of costless adjustment. The proposed state space representation has two blocks: a measurement equation and two transition equations. The measurement equation maps the lagged value of the lumpy variable and the current latent states to the current value of the lumpy variable. Note that such mapping simply captures two-sided (*S*,*s*) rules subject to random opportunities for free adjustments. The transition equations specify the evolution of the two latent states. These can be nonlinear and subject to non-Gaussian disturbances. In addition, we consider an auxiliary block summarizing the evolution of additional observables over time. These could include both micro and macro variables. This framework is able to accommodate as special cases well-known models of (S,s) rules such as random menu cost models (Stokey, 2009; Alvarez et al., 2016; Auclert et al., 2022), models of plant level investment (Baley and Blanco, 2020; Bertola and Caballero, 1994; Caballero et al., 1995), models of individual cash withdrawals (Alvarez and Lippi, 2009; Miller and Orr, 1966; Frenkel and Jovanovic, 1980) and models of labor adjustment (Elsby and Michaels, 2019; Elsby et al., 2019).

Next, we develop state space methods to solve the filtering, smoothing and forecasting problems given our proposed state space representation. The first challenge we face is the curse of dimensionality: the overall state space is of dimension 2^n, given two latent states per observation. Thus, forward filtering and backward smoothing recursions become highly computationally intensive. To overcome this issue, we impose that the shocks entering state transition equations are independent across individuals and time, when conditioned on current and past observable variables.[1] Under this assumption, we can solve for each individual's filtering and smoothing problems *separately*. In other words, instead of a single 2^n-dimensional integral, we now have n two-dimensional integrals, which can be parallelized.

In addition, the presence of (S,s) rules introduces "kinks" in the measurement equation for the lumpy variable as well as the possibility of non-Gaussian state transition errors. This renders the Extended Kalman approach for filtering and smoothing unsuitable for this setup. As an alternative, we use a combination of grid-based filtering and smoothing, and Gauss-Legendre integration. In fact, the former is particularly appealing in an (S,s) economy because the presence of (S,s) rules ensures that the frictionless lumpy variable is *bounded* for each inaction spell. Thus, to approximate the filtered and smoothed densities, we only need to solve the forward filtering and backward smoothing recursions at each Gauss-Legendre node in that bounded interval.

As an empirical application of our framework, we study the effects of monetary policy shocks on inflation in a standard random menu cost model. Prices are set according to a two-sided (S,s) rule subject to an idiosyncratic probability of free adjustment. Moreover, the frictionless price equals a constant markup over the marginal cost. The goal is to construct the impulse response function of inflation to a monetary policy shock. As a first step, we estimate the aggregate input price response from a structural VAR. For identification, we borrow Cesa-Bianchi et al. (2020)'s high-frequency monetary policy surprise series to use as external instrument. Next, we combine data on aggregate input prices with rich micro price data underlying the UK Consumer Price Index to estimate the parameters governing the dynamics of both fric-

[1] Note that this does *not* rule out effects of aggregate shocks on lumpy variables. What we assume is that any *unobserved* aggregate shocks affect lumpy variables only through their effect on the non-lumpy *observable* variables and their effect on the frictionless value of the lumpy variable.

tionless and actual prices together with the cross-sectional distribution of price gaps over time.

The advantage of our framework is that it allows the construction of impulse responses of *any* function of the cross-sectional distribution of lumpy variables given *any* initial distribution of state gaps. We perform four exercises to illustrate the usefulness of our approach. First, we quantify the role of price adjustment frictions in shaping the inflation response to monetary policy by comparing the responses of actual versus frictionless inflation. Second, we decompose the response of inflation at the product, region or household level. Third, we evaluate how the effects of monetary policy vary over time and whether these are state and/or sign dependent depending on the initial distribution of price gaps. Lastly, we provide new inflation forecasts by aggregating individual price trajectory forecasts and test whether they outperform existing inflation forecasting models.

Related Literature. This paper relates to five strands of literature. In terms of methodology, it builds on existing work on state space methods for nonlinear and/or non-Gaussian state space representations (Särkkä, 2013). More specifically, on articles that explore grid-based methods such as Kitagawa (1987) and Kramer and Sorenson (1988). Our paper is closest to Bandeira (2020) that provides a closed-form solution to the filtering and smoothing problems in a model with a single observable variable, which is subject to adjustment frictions and follows a random walk with drift and Gaussian disturbances in its frictionless form. We deviate from Bandeira (2020), however, along two important dimensions: we provide state space methods for a general class of (S,s) economies and study the propagation of aggregate shocks in these settings.

Second, it contributes to a line of research that uses micro and macro data to estimate macroeconomic models (Liu and Plagborg-Møller, 2023; Fernández-Villaverde et al., 2020; Chang et al., 2022). The key difference is that we propose a semi-structural approach, which increases flexibility and helps with parameter estimation. In fact, our state space representation is able to accommodate previous models of microeconomic adjustment frictions by specifying the correct mapping to structural parameters. However, this is not required in order to estimate cross-sectional distributions of latent states, construct impulse response functions and produce forecasts.

Third, the underlying research question is closely related to the literature that looks at the implications of the presence of (S,s) rules for aggregate dynamics. Caballero and Engel (1991) characterize structural and stochastic heterogeneities on one-side dynamic (S,s) economies that ensure the convergence of the aggregate economy to its frictionless equilibrium. Alvarez et al. (2016), Baley and Blanco (2020) and Alvarez and Lippi (2022) provide closed-form expressions for the impulse responses of aggregate variables to a permanent, one-off aggregate shock in terms of sufficient statistics. We view our approach as complementing previous work: we *estimate* the impulse responses of any function of the cross-sectional distribution of lumpy variables to a permanent or transitory aggregate shock and for any initial distribution of price gaps.

Fourth, our empirical application mostly connects to research on heterogeneous effects of monetary policy. Boivin et al. (2009) and Cravino et al. (2020) decompose aggregate inflation responses across sectors and households respectively, using Factor-Augmented Vector Autoregression Models (FAVARs) that do not account for lumpiness. In addition, a number of papers investigate whether the effects of a monetary policy shock vary across time (Primiceri, 2005), depending on the sign of the shock (Long and Summers, 1988; Cover, 1992; Ravn and Sola, 2004) or over the business cycle (Vavra, 2014; Tenreyro and Thwaites, 2016; Angrist et al., 2018). We argue that the key object to explain any heterogeneity in effects is the cross-sectional distribution of price gaps, which we are able to recover empirically.

Finally, there is a vast literature on inflation forecasting, as surveyed by Faust and Wright (2013). Our paper is, to the best of our knowledge, the first to combine micro and macro data and account for microeconomic lumpiness in producing inflation forecasts at any level of disaggregation.

2. A state space representation of a dynamic (S,s) economy

A state space model is a time series model in which a time series, Y_t, is interpreted as the result of a noisy observation of a stochastic process X_t. Its representation usually consists of an initial value for the state variable together with a transition equation, describing the dynamics of the state variables, and a measurement equation, linking the observed variable, the measurement, to the state. We will focus, instead, on the probabilistic state space representation, which uses conditional density functions, rather than equations, to summarize the primitives of the model.

First, we introduce notation. Consider a panel with agents indexed by $i = 1, \ldots, n$ and time indexed by $t = 1, \ldots, T$.[2] $s_{i,t}$ denotes the lumpy measurement for agent i at time t. For each i, there are two state variables, $\mathbf{X}_{i,t} \equiv [s^{\star}_{i,t}, \ell_{i,t}]$, where $s^{\star}_{i,t}$ denotes the frictionless value of $s_{i,t}$ and $\ell_{i,t}$ is an indicator variable that equals one if the agent receives a costless adjustment opportunity.[3] We allow for addional measurement variables, $\tilde{\mathbf{Z}}_t \equiv [\mathbf{z}_{1:n,t}, \mathbf{Z}_t]$, where $\mathbf{z}_{i,t}$ is a vector of agent-specific variables and $\mathbf{Z}_t$ is a vector of aggregate variables. The econometrician, thus, observes $\mathbf{Y}_{1:T} \equiv \{s_{1:n,1:T}, \tilde{\mathbf{Z}}_{1:T}\}$.

[2] To ease exposition, we assume a balanced panel. Note, however, that results still hold otherwise.

[3] The frictionless value of the lumpy variable is the value that the agent would choose in absence of microeconomic adjustment frictions

The probabilistic state space representation of (S,s) economies requires four conditional densities: for the lumpy measurement, $p\left(\mathbf{s}_{1:n,t} \,|\, \mathbf{X}_{1:n,1:t}, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t}; \boldsymbol{\Theta}\right)$; for the non-lumpy measurements, $p\left(\tilde{\mathbf{Z}}_t \,|\, \mathbf{X}_{1:n,1:t-1}, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t-1}; \boldsymbol{\Theta}\right)$; for the states, $p\big(\mathbf{X}_{1:n,t} \,|\, \mathbf{X}_{1:n,1:t-1}, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t}; \boldsymbol{\Theta}\big)$; and for the initial states, $p\left(\mathbf{X}_{1:n,1} \,|\, \mathbf{s}_{1:n,1}, \tilde{\mathbf{Z}}_1; \boldsymbol{\Theta}\right)$. In what follows, we discuss four assumptions of the model, each of which simplifies one of the conditional densities functions above.

A1. Lumpy measurements follow a two-sided (S,s) rule

The first assumption follows naturally from our focus on lumpy economies. To formalize agent behavior, let's define the *inaction function* $d : \mathbb{R}^2 \times \{0,1\} \to \{0,1\}$ as:

$$d(s_{i,t-1}, \mathbf{X}_{i,t}; \boldsymbol{\theta}_{1,i}) = \mathbb{1}\{s_{i,t-1} - s^{\star}_{i,t} \in [L_i, U_i]\}\,(1 - \ell_{i,t}) + \mathbb{1}\{s_{i,t-1} - s^{\star}_{i,t} = c_i\}\,\ell_{i,t}\,, \quad (1)$$

where $\boldsymbol{\theta}_{1,i} \equiv (L_i, c_i, U_i)$ and it is such that $L_i < c_i < U_i$ is a vector of parameters that includes the lower and upper bound of the inaction region (L_i and U_i) and the reset point for the lumpy variable gap (c_i). According to (1) agents choose not to adjust under the two scenarios. First, provided there is no free adjustment, whenever the gap between the lagged lumpy variable and its current frictionless value falls within the inaction region. Second, provided there is no free adjustment, when the gap between the lagged lumpy variable and its current frictionless value is exactly at the reset point.

Thus, $s_{i,t}$ evolves according to the following *individual* conditional probability function:

$$\begin{aligned} p\left(s_{i,t} \,|\, \mathbf{X}_{1:n,1:t}, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t}; \boldsymbol{\Theta}\right) = {} & \delta\left(s_{i,t} - s_{i,t-1}\right) d(s_{i,t-1}, \mathbf{X}_{i,t}; \boldsymbol{\theta}_{1,i}) \\ & + \delta\left(s_{i,t} - (s^{*}_{i,t} + c)\right)\left(1 - d(s_{i,t-1}, \mathbf{X}_{i,t}; \boldsymbol{\theta}_{1,i})\right) \end{aligned} \quad (2)$$

At a given point in time, the individual probability mass function takes either the past value of the lumpy variable with probability given by the inaction indicator or a new value such that the gap is set to the reset point with the opposite probability.

Given that (2) holds for every i and is independent across i's, the conditional density for the lumpy measurement can be written as the product of the individual densities: $p\left(\mathbf{s}_{1:n,t} \,|\, \mathbf{X}_{1:n,1:t}, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t}; \boldsymbol{\Theta}\right) = \prod_{i=1}^{n} p\left(s_{i,t} \,|\, \mathbf{X}_{i,t}, s_{i,t-1}; \boldsymbol{\theta}_{1,i}\right)$.

A2. Conditional on $\mathbf{s}_{1:n,1:t-1}$, non-lumpy measurements $\perp\!\!\!\perp$ $\mathbf{X}_{1:n,1:t-1}$

The second assumption imposes that conditional on the full history of past lumpy measurements, non-lumpy measurements are independent of past states. This assumption has two implications. First, states only affect non-lumpy measurements through the lumpy measurement. Second, lumpy measurements only affect non-lumpy measurements with a one-period delay.

The conditional density for non-lumpy measurements, which we refer to as the *auxiliary model*, can be simplified to: $p\left(\tilde{\mathbf{Z}}_t \mid \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t-1}; \boldsymbol{\theta}_3\right)$

A3. States are Markovian and conditionally independent across i's

The third assumption implies that for every i, only the past state affects the current one and parameters are individual specific. In other words, any common shock must be observable. The conditional density is, once again, a product of the individual ones: $p\left(\mathbf{X}_{1:n,t} \mid \mathbf{X}_{1:n,1:t-1}, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t}; \boldsymbol{\Theta}\right) = \prod_{i=1}^{n} p\left(\mathbf{X}_{i,t} \mid \mathbf{X}_{i,t-1}, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t}; \boldsymbol{\theta}_{2,i}\right)$.

A4. Initial states are conditionally independent across i's

Finally, the individual filtered distribution depends on initial measurements and individual specific parameters and the overall density is the product of the individual ones: $p\left(\mathbf{X}_{1:n,1} \mid \mathbf{s}_{1:n,1}, \tilde{\mathbf{Z}}_1; \boldsymbol{\Theta}\right) = \prod_{i=1}^{n} p\left(\mathbf{X}_{i,1} \mid \mathbf{s}_{1:n,1}, \tilde{\mathbf{Z}}_1; \boldsymbol{\theta}_{0,i}\right)$

A1 is a feature of (S,s) models, while A2 and the Markovian states assumption (A3.a) are implicit in many state space representations. It's mostly the independence assumption, captured by both A3.b and A4, that is strictly our choice. It deals with the curse of dimensionality: the overall state space is of dimension $2n$, which implies evaluating functions at (2 x # grid points)^n nodes to do inference. With independence across i's, each problem can be solved individually, resulting in n problems with a state space dimension of 2 each. Thus, we instead need to evaluate n functions at (2 x # grid points) nodes.

3. State space methods for a dynamic (S,s) economy

Given a data generating process that satisfies the assumptions described in the previous section, we next present the state space methods to estimate the states, $\mathbf{X}_{1:n,1:t-1}$, over time as well as the model parameters. The former requires deriving the filtered probability density, the one-step ahead predicted probability density and the smoothed probability density of states. With the model primitives and the states, we are then able to derive the likelihood function, required to accomplish the latter through Maximum Likelihood or Bayesian Estimation.

Under our four assumptions above, the filtered and one-period ahead predicted probability density for the states are given by the following Lemma:

***Lemma* 1** *Under A1-A4,*

$$p\left(\mathbf{X}_{1:n,t} \,|\, \mathbf{s}_{1:n,1:t}, \tilde{\mathbf{Z}}_{1:t}; \mathbf{\Theta}\right) = \prod_{i=1}^{n} p\left(\mathbf{X}_{i,t} \,|\, \mathbf{s}_{1:n,1:t}, \tilde{\mathbf{Z}}_{1:t}; \mathbf{\Theta}\right) \;,$$
$$p\left(\mathbf{X}_{1:n,t} \,|\, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t}; \mathbf{\Theta}\right) = \prod_{i=1}^{n} p\left(\mathbf{X}_{i,t} \,|\, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t}; \mathbf{\Theta}\right) \;,$$

where,

$$p\left(\mathbf{X}_{i,t} \,|\, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t}; \mathbf{\Theta}\right) = \int p\left(\mathbf{X}_{i,t} \,|\, \mathbf{X}_{i,t-1}, \mathbf{s}_{1:n,1:t}, \tilde{\mathbf{Z}}_{1:t}; \mathbf{\Theta}\right) \times p\left(\mathbf{X}_{i,t-1} \,|\, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t-1}; \mathbf{\Theta}\right) \, d\mathbf{X}_{i,t-1} \tag{3}$$

$$p\left(\mathbf{X}_{i,t} \,|\, \mathbf{s}_{1:n,1:t}, \tilde{\mathbf{Z}}_{1:t}; \mathbf{\Theta}\right) \propto p\left(s_{i,t} \,|\, \mathbf{X}_{i,t}, s_{i,t-1}; \boldsymbol{\theta}_{2,i}\right) \times p\left(\mathbf{X}_{i,t} \,|\, \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t}; \mathbf{\Theta}\right) \tag{4}$$

Proof. See Appendix □

The filtered and one-period ahead predicted densities for the states are equal to the product of the individual densities. The individual predicted density is equal to the product of the conditional density for states (a model primitive) and the individual filtered density at time $t-1$ integrated over all possible states at time $t-1$. The individual filtered density is proportional to the conditional density for the lumpy measurement (a model primitive) and the individual predicted density. To solve for these, take the initial filtered state distribution as a starting point and iterate forward alternating equations (3) and (4) .

Given the one-period ahead predicted density, we derive a similar Lemma for the smoothing problem:

***Lemma* 2** *Under A1-A4,*

$$p\left(\mathbf{X}_{1:n,t} \,|\, \mathbf{s}_{1:n,1:T}, \tilde{\mathbf{Z}}_{1:T}; \mathbf{\Theta}\right) = \prod_{i=1}^{n} p\left(\mathbf{X}_{i,t} \,|\, \mathbf{s}_{1:n,1:T}, \tilde{\mathbf{Z}}_{1:T}; \mathbf{\Theta}\right) \;,$$

where,

$$p\left(\mathbf{X}_{i,t} \,|\, \mathbf{s}_{1:n,1:T}, \tilde{\mathbf{Z}}_{1:T}; \mathbf{\Theta}\right) = p\left(\mathbf{X}_{i,t} \,|\, \mathbf{s}_{1:n,1:t}, \tilde{\mathbf{Z}}_{1:t}; \mathbf{\Theta}\right) \times$$
$$\int \frac{p\left(\mathbf{X}_{i,t} \,|\, \mathbf{X}_{i,t-1}, \mathbf{s}_{1:n,1:t}, \tilde{\mathbf{Z}}_{1:t}; \mathbf{\Theta}\right) p\left(\mathbf{X}_{i,t+1} \,|\, \mathbf{s}_{1:n,1:T}, \tilde{\mathbf{Z}}_{1:T}; \mathbf{\Theta}\right)}{p\left(\mathbf{X}_{i,t+1} \,|\, \mathbf{s}_{1:n,1:t}, \tilde{\mathbf{Z}}_{1:t+1}; \mathbf{\Theta}\right)} \, d\mathbf{X}_{i,t+1}$$

Proof. See Appendix □

The individual smoothed density is equal to the product of two elements: the individual smoothed distribution at time t and the conditional density for states times next period's smoothed density over the one-period ahead predicted density integrated over individual states at time $t+1$. The solution now requires backward iteration and the starting point is the last filtered state.

We are now ready to derive the likelihood function as follows:

***Lemma* 3** *Under A1-A4,*

$$p\left(\mathbf{Y}_{1:T} \mid \boldsymbol{\Theta}\right) = p\left(\mathbf{Y}_1 \mid \boldsymbol{\Theta}\right) \text{ x } \prod_{t=2}^{T} p\left(\tilde{\mathbf{Z}}_t \mid \mathbf{s}_{1:n,1:t-1}, \tilde{\mathbf{Z}}_{1:t-1}; \boldsymbol{\theta}_3\right)$$

$$\text{x } \prod_{i=1}^{n}\prod_{t=2}^{T} \int p\left(s_{i,t} \mid \mathbf{X}_{i,t}, s_{i,t-1}; \boldsymbol{\theta}_{1,i}\right) \, p\left(\mathbf{X}_{i,t} \mid \mathbf{s}_{i:n,t-1}, \tilde{\mathbf{Z}}_{1:t-1}; \boldsymbol{\Theta}\right) d\mathbf{X}_{i,t}$$

Proof. See Appendix □

The likelihood is the product of the initial conditions, the auxiliary model and the likelihood for lumpy measurements i.e. the product of the conditional density for lumpy measurement and the one-step ahead predicted density across time and across individuals. A key feature of the likelihood is parameter separability: $\boldsymbol{\theta}_3$ can be estimated separately from $\{\boldsymbol{\theta}_i\}_{i=1}^{n}$. Moreover, if we make the assumption that $\boldsymbol{\theta}_i$'s are common within groups, we shall evaluate likelihood at the group level.

Grid Based Methods. Developing state space methods for lumpy economies poses one key challenge: the kinks. This non-linearity, together with the possibility of non-Gaussian state transition errors, rules out closed-form solutions to the forward filtering and backward smoothing recursions defined above. In other words, widely used techniques such as the Kalman filter and its extended family are no longer suitable. At the same time, the (S,s) rules that generate the kinks lead the way towards the use of grid based approximation techniques.

Grid based methods discretize the state space and solve the recursions only at a selected number of points. The hurdle is point selection, in particular, keeping the number of points as small as possible. There are two state in the state space representation: one has only a two-point support but the second is continuous. Given the (S,s) rule, however, $s_{i,t}$ is continuous but bounded. Bounds allow quadrature methods, which are known to provide good approximations with a small number of grid points. In particular, we use Gauss-Legendre nodes.

4. Monetary Policy and Inflation in a Random Menu Cost Economy

As an application of our framework we study the propagation of monetary policy shocks in a random menu cost economy. In this application the lumpy variable is taken to be (log) prices of different products observed in micro price data whose measurement equation is given by (2) and the arrival of costless adjustment opportunities is given by (**??**). For (**??**) and (**??**) we assume the following,

Frictionless prices. Each price observed in the micro data is set by a single product firm that has the following production function,

$$Y_{i,t} = A_{i,t} K_{i,t}^{\alpha} N_{i,t}^{1-\alpha} \tag{5}$$

where the two production inputs are capital ($K_{i,t}$) and labor ($N_{i,t}$) purchased by each firm at prices R_t and W_t in perfectly competitive markets.[4] The idiosyncratic total factor productivity (TFP) component is assumed to evolve according to,

$$a_{i,t} = \rho a_{i,t-1} + \varepsilon_{i,t} \tag{6}$$

where $a_{i,t} \equiv \log A_{i,t}$ and $\varepsilon_{i,t}$ is an idiosyncratic TFP shock distributed according to some probability density function, $p(\varepsilon_{i,t})$, and *i.i.d.* across i and t.[5] Products are sold under monopolistic competition and each firm faces an isoelastic demand curve. Under these assumptions, frictionless prices (in logs) evolve according to,

$$p^{\star}_{i,t} = \beta_0 + \rho p^{\star}_{i,t-1} + \beta_{r,0}\, r_t + \beta_{r,1}\, r_{t-1} + \beta_{w,0}\, w_t + \beta_{w,1}\, w_{t-1} - \varepsilon_{i,t} \tag{7}$$

where $r_t \equiv \log R_t$, $w_t \equiv \log W_t$ and the β's are functions of the structural parameters (in this case α, ρ and the elasticity of demand ε).

Input prices. In this application there are no other non-lumpy observables at the firm level that affect frictionless prices so we only need to specify an auxiliary transition equation for aggregate input prices. It is assumed that (log) of input prices follows a Vector Autoregression (VAR) augmented with a function of the cross-sectional distribution of (log) prices as follows,

[4] For simplicity we assume that capital fully depreciates after one period. Note that (5) could be easily generalized to include more inputs.

[5] In our framework, TFP shocks can be distributed according to a normal distribution (Golosov and Lucas, 2007), a mixture of normal distribution (Karadi and Reiff, 2019) or a fat-tailed distribution (Gertler and Leahy, 2008; Midrigan, 2011).

$$\mathbf{Z}_t = \mathbf{A}_0 + \sum_{j=1}^{p} \mathbf{A}_j \mathbf{Z}_{t-j} + \sum_{j=1}^{p} \mathbf{C}_j h(p_{1:n,t-j-1:t-j}) + \mathbf{u}_t \tag{8}$$

where $\mathbf{Z}_t = [w_t, r_t]'$ is a vector of input prices, $h(\cdot)$ is a known function of the cross-sectional distribution of prices (*e.g.* the aggregate price index) and $\mathbf{u}_t$ is *i.i.d.* white noise with variance $\mathbf{\Sigma}_u$. As commonly done in the SVAR literature, we assume that there exists a rotation $\mathbf{B}_0$ such that $\mathbf{B}_0\mathbf{u}_t = \varepsilon_t$ where ε_t is a vector of structural shocks that contains the monetary policy shock. Finally, monetary policy shocks are identified using the high-frequency surprises from Cesa-Bianchi et al. (2020) as an external instrument.

Data and estimation. We use the methods presented in section 3, rich micro price data underlying the UK CPI and aggregate data on input prices to: (*i*) estimate all the parameters in (2), (**??**), (7) and (8), (*ii*) estimate the cross-sectional distributions of latent states over time and (*iii*) compute the impulse response functions of inflation at different disaggregation levels at different points in time.

References

Alvarez, Fernando and Francesco Lippi, "Financial Innovation and the Transactions Demand for Cash," *Econometrica*, March 2009, *77* (2), 363–402.

_ **and** _ , "The Analytic Theory of a Monetary Shock," *Econometrica*, July 2022, *90* (4), 1665–1680.

_ **, Hervé Le Bihan, and Francesco Lippi**, "The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach," *American Economic Review*, October 2016, *106* (10), 2817 – 2851.

Angrist, Joshua D., Òscar Jordà, and Guido M. Kuersteiner, "Semiparametric Estimates of Monetary Policy Effects: String Theory Revisited," *Journal of Business and Economic Statistics*, 2018, *36* (3), 371–387.

Auclert, Adrien, Rodolfo Rigato, Matthew Rognlie, and Ludwig Straub, "New Pricing Models, Same Old Phillips Curves?," April 2022. Manuscript.

Baley, Isaac and Andrés Blanco, "Aggregate Dynamics in Lumpy Economies," *Econometrica*, May 2020, *89* (3), 1235–1264.

Bandeira, Miguel, "Essays in Macroeconometrics." PhD dissertation, London School of Economics and Political Science July 2020.

Bertola, Giuseppe and Ricardo J. Caballero, "Irreversibility and Aggregate Investment," *Review of Economic Studies*, 1994, *61*, 223–246.

Boivin, Jean, Marc P. Giannoni, and Ilian Mihov, "Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data," *The American Economic Review*, March 2009, *99* (1), 350–384.

Caballero, Ricardo, Eduardo Engel, and John Haltiwanger, "Plant-Level Adjustment and Aggregate Investment Dynamics," *Brookings Papers on Economic Activity*, 1995, *2.*

Caballero, Ricardo J. and Eduardo M. R. A. Engel, "Dynamic (S,s) Economies," *Econometrica*, November 1991, *59* (6), 1659 – 1686.

Cesa-Bianchi, Ambrogio, Gregory Thwaites, and Alejandro Vicondoa, "Monetary policy transmission in the United Kingdom: A high frequency identification approach," *European Economic Review*, October 2020, *123.* Working Paper.

Chang, Minsu, Xiaohing Chen, and Frank Schorfheide, "Heterogeneity and Aggregate Fluctuations," October 2022. Working Paper.

Cover, James Peery, "Asymmetric Effects of Positive and Negative Money-Supply Shocks," *The Quarterly Journal of Economics*, November 1992, *107* (4), 1261–1282.

Cravino, Javier, Ting Lan, and Andrei A. Levchenko, "Price stickiness along the income distribution and the effects of monetary policy," *Journal of Monetary Economics*, April 2020, *110*, 19–32.

Elsby, Michael W. L. and Ryan Michaels, "Fixed Adjustment Costs and Aggregate Fluctuations," *Journal of Monetary Economics*, January 2019, *101*, 128 – 147.

_ , _ , and David Ratner, "The aggregate effects of labor market frictions," *Quantitative Economics*, 2019, *10*, 803–852.

Faust, Jon and Jonathan H. Wright, "Inflation Forecasting," in Graham Elliot and Allan Timmerman, eds., *Handbook of Economic Forecasting*, Vol. 2A, Elsevier, 2013, pp. 2–56.

Fernández-Villaverde, Jesús, Samuel Hurtado, and Galo Nuno, "Financial Frictions and the Wealth Distribution," 2020. Banco de España, Documentos de Trabajo, n 2013.

Frenkel, Jacob A. and Boyan Jovanovic, "On Transaction and Precautionary Demand for Money," *The Quarterly Journal of Economics*, August 1980, *95* (1), 25–43.

Gertler, Mark and John Leahy, "A Phillips Curve with an Ss Foundation," *Journal of Political Economy*, July 2008, *116* (3), 533 – 572.

Golosov, Mikhail and Robert E. Lucas, "Menu Costs and Phillips Curves," *Journal of Political Economy*, April 2007, *115* (2), 171–199.

Karadi, Peter and Adam Reiff, "Menu Costs, Aggregate Fluctuations, and Large Shocks," *American Economic Journal: Macroeconomics*, July 2019, *11* (3), 111–146.

Kitagawa, Genshiro, "Non-Gaussian State-Space Modeling of Nonstationary Time Series," *Journal of the American Statistical Association*, December 1987, *82* (400), 1032–1041.

Kramer, Stuart C. and Harold W. Sorenson, "Recursive bayesian estimation using piecewise constant approximations," *Automatica*, November 1988, *24* (6), 789–801.

Liu, Laura and Mikkel Plagborg-Møller, "Full-Information Estimation of Heterogeneous Agent Models Using Macro and Micro Data," *Quantitative Economics*, January 2023, *14* (1), 1–35.

Long, J. Bradford De and Lawrence H. Summers, "How Does Macroeconomic Policy Affect Output?," *Brookings Papers on Economic Activity*, 1988, *2*, 433–494.

Midrigan, Virgiliu, "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations," *Econometrica*, July 2011, *79* (4), 1139–1180.

Miller, Merton H. and Daniel Orr, "A Model of the Demand for Money by Firms," *The Quarterly Journal of Economics*, August 1966, *80* (3), 413–435.

Primiceri, Giogio E., "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*, July 2005, *72* (3), 821–852.

Ravn, Morten O. and Martin Sola, "Asymmetric Effects of Monetary Policy in the United States," *Federal Reserve Bank of St. Louis Review*, September/October 2004, *86* (5), 41–60.

Särkkä, Simo, *Bayesian Filtering and Smoothing* Institute of Mathematical Statistics Textbooks, first ed., Cambridge University Press, 2013.

Stokey, Nancy L., *The Economics of Inaction: Stochastic Control Models with Fixed Costs*, first ed., Princeton University Press, 2009.

Tenreyro, Silvana and Gregory Thwaites, "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, October 2016, *8* (4), 43–74.

Vavra, Joseph, "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *The Quarterly Journal of Economics*, February 2014, *129* (1), 215–258.