Economics Bulletin

Volume 46, Issue 2

AI-tocracy dynamics

Orlando Gomes
Lisbon Accounting and Business School (ISCAL-IPL) and CEFAGE-ISCAL research center

Abstract

AI-tocracy is the designation used to characterize a political regime that is autocratic and that resorts to AI-based methods of control and repression to guarantee its entrenchment in power. This study develops an original dynamic model of AI-tocracy, in which political control is pursued via investment in digital surveillance. The sketched model allows for different degrees of authoritarianism and for different levels of efficacy of surveillance. Numerical results indicate that, in the steady state, deeply autocratic regimes can coexist with a fast expansion of the AI frontier, as long as the efficacy of surveillance is high (thus creating prosperous AI-tocracies). If the government is incapable of efficaciously extracting data from socio-economic events and interactions, then deeper authoritarianism will be accompanied by the decline of AI technology and by a fall in output and welfare.

The author gratefully acknowledges the comments of an anonymous referee and of the journal´s Editor (the usual disclaimer applies).The author is also pleased to acknowledge the financial support from Fundação para a Ciência e a Tecnologia (Portugal) under the Project UID/04007/2025.
Citation: Orlando Gomes, (2026) "AI-tocracy dynamics", *Economics Bulletin*, Volume 46, Issue 2, pages 604-612
Contact: Orlando Gomes - omgomes@iscal.ipl.pt.
Submitted: February 18, 2026. **Published:** June 30, 2026.

1. Introduction: The AI Era of Political Control

Recently, a new term – AI-tocracy – entered the lexicon of economists and political scientists (Beraja *et al*., 2023*a*; Bai *et al*., 2023; Yang, 2024). The term suggests a close relationship between political control in autocratic regimes and artificial intelligence (AI) technology. AI enables large-scale monitoring and an effective prediction of human behavior, allowing the autocrat's toolkit of surveillance, control, manipulation, and repression, to be significantly strengthened.

The AI-tocracy argument unfolds as follows: because authoritarian regimes derive political power from AI, they have an incentive to provide institutional and financial support to firms endowed with the capacity to develop the desired tech solutions. With reinforced political protection and increased investment, the regime-backed firms will be able to expand the technology frontier, for the intended political purposes and, also, through innovation spillovers, for civil and commercial uses.

A process of mutual reinforcement is triggered: because the government desires to employ digital technology to exert social control, it will entitle the economy with the resources needed to expand the AI innovation frontier (including data collected through digital surveillance). As AI becomes more powerful, also the means of control become more effective, consolidating the entrenchment of the regime. Under this logic, the notion of AI-tocracy can be associated with a political and socio-economic equilibrium where autocracy and AI innovation become two faces of a same coin.

In the sections that follow, an intertemporal optimization problem is devised to address the AI-tocracy equilibrium. The goal is to identify the conditions in which the best AI outcome coincides with autocracy taken to its fullest extent. Although relatively simple, the analytical model involves nonlinearities that make it unfeasible to undertake a general assessment of the respective optimality conditions. Numerical examples are explored, revealing that the outcome is contingent on parameter values and, thus, on the configuration of some relevant functions, as it is potentially the case of the data collection function and of the production function for AI.

The remainder of the manuscript puts together the pieces required to formulate the autocrat's intertemporal optimization problem (sections 2 to 4), numerically evaluates the steady state (section 5), and briefly discusses results (section 6).

2. Digital Surveillance and Data Availability

Data is a by-product of economic activity. As market transactions unfold, a massive volume of digitized information is created. Modern data science technologies make it possible to process and analyze the generated digital data sets with impressive efficiency and accuracy, what constitutes a major advantage for companies in attracting customers, managing inventories, customizing products, and forecasting sales. Following the literature on the economics of data (Farboodi *et al*., 2019; Farboodi and Veldkamp, 2023; Veldkamp and Chung, 2024; Baley and Veldkamp, 2025), the flow of data in the economy at time t, $n(t) \geq 0$, is modelled as a contemporary linear function of aggregate output, $y(t) \geq 0$, with parameter $z > 0$ representing data savviness, i.e., the ability of firms in extracting valuable data from production and trade: $n(t) = zy(t)$.

Social monitoring systems, which are common in most authoritarian regimes, can be a powerful tool in extracting additional data, beyond the one strictly enabled by economic transactions (Beraja *et al.,* 2023*b*). Notwithstanding, some authors (Yang and Roberts, 2023; Woods, 2025; Yang, 2025) argue in favor of a nonlinear relationship between digital surveillance and the corresponding data collection outcome. The argument is that excessive control and repression may induce a change in people's behavior (self-censorship and preference falsification) that severely constrains the ability of public authorities in gathering accurate and unbiased data.

The above reasoning compels us to modify the data flow function, by introducing a surveillance component. Let $G(t) \geq 0$ denote the government's spending on social monitoring. At low levels of $G(t)$, this spending effectively promotes access to additional quality data. Beyond a given threshold of surveillance intensity, though, the data extraction capabilities progressively degrade and the social control effort by the authorities becomes counter-productive.

Analytically, we choose to model the characterized logic under the following data flow function:

$$n(t) = zM^{1-\left(\frac{G(t)}{a}-1\right)^2} y(t), \qquad M > 1, a > 0 \tag{1}$$

In Eq.(1), parameter M stands for the maximum contribution given by surveillance to the process of data extraction, which is attained at a monitoring spending level $G(t) = a$. Observe that $n(t)$ is a hump-shaped function of $G(t)$, with $n(t) = zy(t)$ for $G(t) = 0$, $n(t) = zMy(t)$ for $G(t) = a$, and $n(t) = 0$ for $G(t) \rightarrow \infty$ (in a hypothetical scenario of infinite spending in surveillance, the State would cease to exist and no data would be collected).

Data can be accumulated over time. Let $\Omega(t) \geq 0$ represent the stock of data and take $\zeta \in (0,1)$ as the rate of data depreciation. Assuming time flows continuously, the process of data accumulation is given by the following differential equation,

$$\frac{d\Omega(t)}{dt} = n(t) - \zeta\Omega(t), \qquad \Omega(0) \text{ given} \tag{2}$$

3. Sketching a Production Function for AI

In its essence, AI technology is fundamentally the outcome of a process of data training (Karpa *et al.*, 2022; Klarl *et al.,* 2023). Defining $A(t) \geq 0$ as an AI index, one can discuss which functional form best characterizes the relationship between the stock of data and AI. We define a continuous and differentiable production function $A(t) = f[\Omega(t)]$ that obeys some trivial properties; namely, $f(0) = 0$ and $f' > 0$. The sign of the second derivative of function f is not as straightforward to establish as the sign of the first derivative; although logically AI should benefit from a larger stock of data, the ways in which the corresponding growth accelerates or decelerates with data accumulation depends on a series of factors that can be integrated into the production function and that eventually allow to distinguish different phases in the relationship between data availability and the ability to enhance the capabilities of digital technologies.

Three effects are accounted for in assembling the AI production function. The first is a data scale effect such that AI directly benefits from the amount of available data (associated marginal returns are allowed to be diminishing, constant, and increasing); the second is a low-data penalty, according to which low data levels prevent AI technology to take off; and the third is a saturation effect suggesting that, after a given threshold, the addition of more data is incapable of significantly expanding the state of technology. These three effects appear, in the listed order, in the production function displayed in Eq.(3),

$$A(t) = A_0\Omega(t)^{\omega}e^{-\beta/\Omega(t)}\left(1 - e^{-\gamma\Omega(t)}\right), \qquad A_0, \omega, \beta, \gamma > 0 \tag{3}$$

Eq.(3) is a nonlinear production function capturing the various ways in which the availability of data impacts the development of AI technology. This is a general-purpose technology, which can acquire many forms and serve distinct economic and social goals. E.g., it can materialize into a sales prediction tool for firms or a facial recognition mechanism for social control. The various possible uses of AI are captured in our model by assuming that the technology is an input in the production of goods and a vehicle employed by the political regime to reinforce social control.

On the economic side, we consider a Cobb-Douglas production function with a constant and normalized to 1 labor force,

$$y(t) = A(t)k(t)^{\alpha}, \qquad \alpha \in (0,1) \tag{4}$$

In Eq.(4), variable $k(t) \geq 0$ represents physical capital, and parameter α stands for the output-capital elasticity. AI technology can be interpreted, in the context of the production function, as a productivity measure. Capital accumulation follows a trivial dynamic rule,

$$\frac{dk(t)}{dt} = y(t) - c(t) - T(t) - \delta k(t), \qquad \delta \in (0,1) \tag{5}$$

In Eq.(5), δ is the capital depreciation rate, variable $c(t) \geq 0$ represents the level of private consumption, and $T(t) \geq 0$ is the tax level. In this economy, taxes serve a single purpose, which is to finance the government's effort in extracting data through social surveillance. Assuming a balanced budget, $T(t) = G(t)$.

Concerning political control, one may assume a parameter $\Lambda_0 > 0$ to symbolize the level of control that public authorities exert through traditional means of coercion and oppression. This level of control can be magnified through the implementation of AI technologies. A function that captures this combined impact is the following:

$$\Lambda(t) = \ln\left[e^{\Lambda_0} + \lambda A(t)\right], \qquad \lambda > 0 \tag{6}$$

Observe, in Eq.(6), that if $A(t) = 0$ then $\Lambda(t) = \Lambda_0$, and that $\frac{d\Lambda(t)}{dA(t)} > 0$ and $\frac{d^2\Lambda(t)}{dA(t)^2} < 0$ (control and repression benefit from digital technologies, but there are associated diminishing returns).

4. Different Shades of Autocracy: the Ruler's Choices and Constraints

To classify political regimes, one may adopt an autocracy scale (Tohmé *et al.,* 2022; Stockemer, 2025). In one extremity of this scale, we place a fully autocratic government, whose exclusive concern is to preserve power. In the opposed end, we find a government solely worried with social welfare and the maximization of consumers' utility. In between, a wide range of possibilities exist, signifying different possible shades of authoritarianism.

A ruler's objective function that translates the idea of an autocracy scale as discussed above can be specified in the following terms:

$$V(t) = (1-\psi)\Lambda(t) + \psi U[c(t)] \tag{7}$$

In Eq.(7), $\psi \in [0,1]$ is a parameter representative of the extent in which the social planner is concerned with consumers' utility. Utility function $U[c(t)]$ obeys trivial conditions of continuity, differentiability, and concavity. For the sake of simplicity, henceforth we consider an elementary logarithmic function for utility, i.e., $U[c(t)] = \ln[c(t)]$. Observe that if $\psi = 0$ then the only concern of the authority is political control; if $\psi = 1$, then the focus is exclusively on consumption utility (we could make this second scenario correspond to a well-functioning democracy, where the government's objective function coincides with the felicity function of the private economy).

Parameter ψ may represent the degree of benevolence of the ruler, but it can also translate the constraints those in power face, e.g., from a legal or constitutional point of view, and that restrain them in their temptation to adopt increasingly dictatorial practices. Clearly, a nexus exists between $\Lambda(t)$ and ψ (additional political control will allow to tear down the barriers that constrain the government's autocratic impulses); however, for the sake of simplicity, in what follows, we maintain these as two separate entities.

With the above logic in mind, an optimal planning problem is assembled. In this problem, the political leader maximizes the present value of the intertemporal stream of objective functions (7), given the value of ψ. The control variables in the problem are $G(t)$ and $c(t)$, i.e., the decision-maker selects the trajectories of public investment in surveillance and private consumption that maximize $V_0 = \int_0^{+\infty} e^{-\rho t} V(t)dt$, with $\rho \geq 0$ the rate of time preference.

The optimal control problem is subject to two dynamic constraints, (2) and (5), given the definitions of data flow, AI, output, and political control in equations (1), (3), (4), and (6). The maximization problem can be solved by resorting to conventional optimal control techniques (see appendix); however, given the nonlinear nature of many of the characterized relationships, no intelligible generic results can be derived. In the next section, a numerical example is explored, to identify the eventual formation of an AI-tocracy equilibrium.

5. A Numerical Investigation of the Authoritarianism Innovation Nexus

The designed model involves fourteen parameters. To pursue a tractable numerical example, twelve of these parameters (all except ψ and a) are assumed constant, taking the values displayed in Table 1. With the selected parameterization, one computes the steady state for the most meaningful variables in the model. In this simulation exercise, parameter ψ crosses the spectrum of autocracy levels (i.e., it assumes values in the range 0 to 1), and distinct scenarios are considered regarding the efficacy of digital surveillance investment. Tables 2 to 4 display the relevant results for three specifications of parameter a: $a = 0.20$, $a = 0.35$, and $a = 0.50$.[1]

Parameter	**Value**
Data depreciation rate	$\zeta = 0.10$
Capital depreciation rate	$\delta = 0.08$
Data savviness	$z = 0.30$
Maximum multiplier on data flow from surveillance	$M = 2.00$
Technology parameter	$A_0 = 1.00$
Elasticity of AI with respect to the data stock	$\omega = 0.30$
Low-data penalty parameter	$\beta = 0.50$
Saturation parameter of AI at high data levels	$\gamma = 0.80$
Output-capital elasticity	$\alpha = 0.33$
Rate of time preference	$\rho = 0.03$
Non-AI political control	$\Lambda_0 = 1.00$
Effectiveness of AI in strengthening political control	$\lambda = 2.00$

Table 1 – Parameter values

	$\boldsymbol{k^*}$	$\boldsymbol{\Omega^*}$	$\boldsymbol{c^*}$	$\boldsymbol{G^*}$	$\boldsymbol{A^*}$	$\boldsymbol{y^*}$	$\boldsymbol{n^*}$	$\boldsymbol{\Lambda^*}$
$\boldsymbol{\psi = 0.00}$	5.887	3.421	0.000	0.418	1.128	1.888	0.342	1.444
$\boldsymbol{\psi = 0.25}$	5.632	3.215	0.681	0.392	1.087	1.811	0.321	1.419
$\boldsymbol{\psi = 0.50}$	5.241	2.902	1.106	0.354	1.026	1.693	0.290	1.382
$\boldsymbol{\psi = 0.75}$	4.683	2.482	1.382	0.303	0.937	1.524	0.248	1.327
$\boldsymbol{\psi = 1.00}$	3.962	1.966	1.548	0.240	0.821	1.302	0.197	1.252

Table 2 – Steady state equilibria ($\boldsymbol{a = 0.20}$).

	$\boldsymbol{k^*}$	$\boldsymbol{\Omega^*}$	$\boldsymbol{c^*}$	$\boldsymbol{G^*}$	$\boldsymbol{A^*}$	$\boldsymbol{y^*}$	$\boldsymbol{n^*}$	$\boldsymbol{\Lambda^*}$
$\boldsymbol{\psi = 0.00}$	5.321	2.801	0.000	0.582	1.038	1.768	0.280	1.382
$\boldsymbol{\psi = 0.25}$	5.632	3.012	0.673	0.523	1.064	1.811	0.301	1.404
$\boldsymbol{\psi = 0.50}$	5.742	3.102	1.098	0.484	1.072	1.843	0.310	1.408
$\boldsymbol{\psi = 0.75}$	5.542	2.921	1.372	0.412	1.042	1.780	0.292	1.404
$\boldsymbol{\psi = 1.00}$	4.842	2.521	1.520	0.323	0.921	1.556	0.252	1.337

Table 3 – Steady state equilibria ($\boldsymbol{a = 0.35}$).

[1] The values in the tables were computed with the assistance of the AI platform DeepSeek-V3.1 (www.deepseek.com).

	k^*	Ω^*	c^*	G^*	A^*	y^*	n^*	Λ^*
$\psi = 0.00$	4.842	2.102	0.000	0.723	0.921	1.556	0.210	1.337
$\psi = 0.25$	5.321	2.521	0.665	0.623	0.982	1.662	0.252	1.373
$\psi = 0.50$	5.632	2.821	1.088	0.562	1.008	1.768	0.282	1.388
$\psi = 0.75$	5.742	2.902	1.362	0.484	1.026	1.811	0.290	1.382
$\psi = 1.00$	5.032	2.521	1.480	0.392	0.921	1.588	0.252	1.337

Table 4 – Steady state equilibria ($a = 0.50$).

Recall that the higher the value of a, the larger is the amount of surveillance expenditure (G) required to attain maximum potential data extraction (M). Hence, a is a measure of government inefficiency; the higher its value, the less effective public authorities are in promptly collecting incremental levels of data. This will reflect negatively on the value of the AI index and, indirectly, on output and the state capacity to exert political control. The negative impact of a larger a is visible in the tables when assessing the values of steady state consumption: consumption grows with the benevolence of the government (higher ψ) but results are more significant for more effective surveillance investment (lower a).

For the other variables, results are not as straightforward, being possible to identify the following pattern:

- Efficient surveillance (a relatively low value of a, i.e., $a = 0.2$) favors authoritarianism. In this case, the government would be better off with a low value of ψ. A low ψ does not only favor political control (Λ^*) but it is also beneficial for data extraction, data accumulation, capital accumulation, output, and technology (AI).
- Less efficient surveillance (higher values of a, i.e., $a = 0.35; a = 0.5$) leads to a different result: every indicator (except consumption) points to a hump-shaped relation with the degree of authoritarianism. Moderately benevolent autocrats benefit from more control, at the same time they allow for more data accumulation, capital accumulation, output, and technological development.

Remember that the central relationship one intends to highlight is the one between the extent of autocracy and AI technology. At this level, results show that as the efficiency in surveillance falls, autocracy loses its capability of promoting the best possible AI outcome, and a less authoritarian regime is capable of delivering relatively better technological outcomes. Hence, AI-tocracy requires special conditions to thrive; one of these conditions, explored here, is how effective the government is in collecting data from citizens and socio-economic interactions among them.

Fig.1 displays, for different values of a (those in tables 2 to 4, and three others) the AI-tocracy relationship. The graphic explicitly indicates that a fully autocratic government correlates with high AI performance at low values of a; however, this relation is progressively subverted as the surveillance capabilities of the government decline and a increases. For relatively high values of a (larger than $a = 0.65$ in the example), the best AI outcome is obtained for a fully democratic welfare-enhancing government ($\psi = 1$). When comparing the various lines in the graphic, one realizes that the absolute best AI outcome is attainable for a full dictatorship regime when a is the lowest possible.

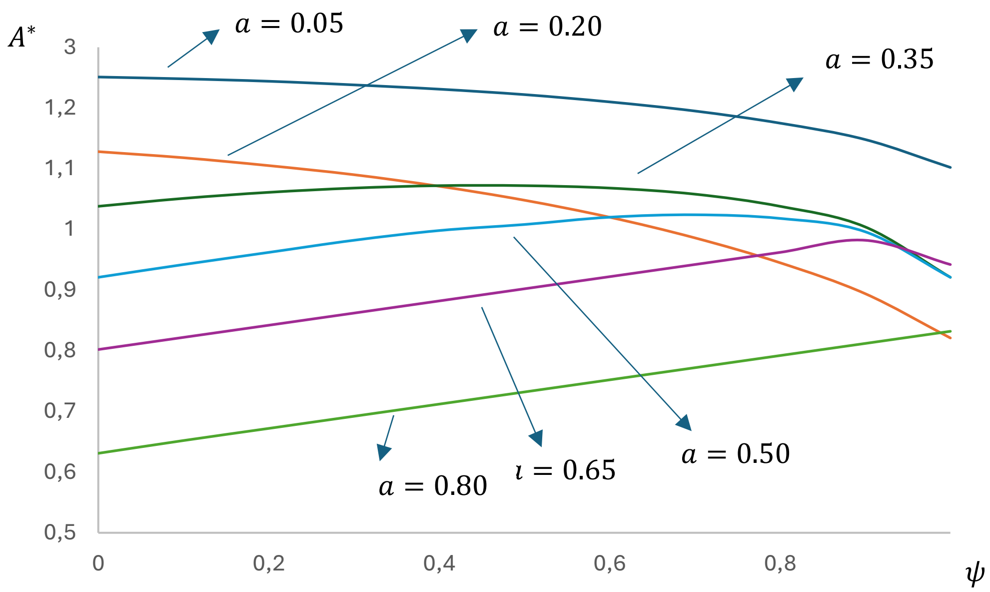

Fig.1 – Authoritarianism and AI innovation.

6. Epilogue: Dictator's Trade-offs and Dilemmas

In the conceived model, the political ruler chooses an optimal trajectory for the level of investment in digital surveillance. Both too low or too high investment levels can compromise data collection and lead to an AI deadlock that is detrimental for the economy and also for the government's desire to enhance political dominance via digital means. The derived optimal surveillance expenditure is not independent from the socio-economic context. It is contingent on the type of regime (i.e., the extent to which political control and social welfare are weighed against each other in the ruler's objective function), and on the efficacy of the government spending on surveillance.

Solving the model and simulating its steady state resorting to a sensible parameterization, one concludes that the link between the degree of authoritarianism and the AI tech capabilities of the economy corresponds to a nonlinear association. Specifically, when the government is not too effective in collecting data through surveillance, the exercise of authority is not conducive to good AI results. In this case, democratic regimes are the ones most capable of the technological synergies required for AI growth. However, if the political regime can set up an effective surveillance campaign, which collects large volumes of data with a relatively small investment, the AI-tocracy equilibrium can be accomplished.

In the AI-tocracy equilibrium there is a symbiotic relation between political control and AI, with big data serving as the middleman: a strong state-controlled investment in surveillance allows for a multiplier effect on the availability of data. The abundance of data warrants to tech firms the possibility of expanding the AI frontier, thus leading to increased productivity in the production of goods and services, and to the development of the tools required by the regime to secure political control. The AI-tocracy equilibrium tends to be self-reinforcing.

References

Bai, L.; L. Wei; and L. Xue (2023) "Endogenous AI-tocracy" SSRN working paper nº 4636843.

Baley, I. and L.L. Veldkamp (2025) *The Data Economy. Tools and Applications*, Princeton, NJ: Princeton University Press.

Beraja, M.; A. Kao; D.Y. Yang; and N. Yuchtman (2023*a*) "AI-tocracy" *Quarterly Journal of Economics*, **138**(3), 1349-1402.

Beraja, M.; D.Y. Yang; and N. Yuchtman (2023*b*) "Data-intensive Innovation and the State: Evidence from AI Firms in China" *Review of Economic Studies*, **90**(4), 1701-1723.

Farboodi, M. and L. Veldkamp (2023) "Data and Markets" *Annual Review of Economics*, **15**(1), 23-40.

Farboodi, M.; R. Mihet; T. Philippon; and L. Veldkamp (2019) "Big Data and Firm Dynamics" *American Economic Association Papers and Proceedings*, **109**, 38-42.

Karpa, D.; T. Klarl; and M. Rochlitz (2022) "Artificial Intelligence, Surveillance, and Big Data" In L. Hornuf (ed.), *Diginomics Research Perspectives. Advanced Studies in Diginomics and Digitalization*, ch. **8**, pp. 145-172.

Klarl, T.; D.F. Karpa; M.E. Leusin; and M. Rochlitz (2023) "Authoritarian Surveillance, Innovation and Growth" SSRN working paper nº 4594849.

Stockemer, D. (2025) "Is the US Moving Toward Autocracy? A Critical Assessment" *Politics & Policy*, **53**(3), e70032.

Tohmé, F.; M.A. Caraballo; and C. Dabús (2022) "Instability, Political Regimes and Economic Growth. A Theoretical Framework" *Metroeconomica*, **73**(1), 291-317.

Veldkamp, L. and C. Chung (2024) "Data and the Aggregate Economy" *Journal of Economic Literature*, **62**(2), 458-484.

Woods, D. (2025) "AI as a Tool for Surveillance: China's Concave Trilemma" *Journal of Chinese Political Science*, 1-35.

Yang, D.Y. (2024) "China: Autocracy 2.0" NBER working paper nº w32993.

Yang, E. (2025) "The Limits of AI for Authoritarian Control" Working paper, Department of Political Science, Purdue University.

Yang, E. and M.E. Roberts (2023). "The Authoritarian Data Problem." *Journal of Democracy*, **34**(4), 141-150.

Appendix – Solution of the Optimal Control Problem

Given the specification of the optimization problem, one can write the corresponding current-value Hamiltonian function,

$$\begin{aligned} H[k(t),\Omega(t),c(t),G(t),p(t),q(t)] &= (1-\psi)\Lambda(t)+\psi U[c(t)]+p(t)[y(t)-c(t)-G(t) \\ &\quad -\delta k(t)]+q(t)[n(t)-\zeta\Omega(t)] \end{aligned} \tag{8}$$

Variables $p(t)$ and $q(t)$ are the shadow-prices associated with each state variable (stocks of capital and data, respectively). First-order conditions are:

$$\frac{\partial H}{\partial c} = 0 \Rightarrow \psi U_c = p(t) \tag{9}$$

$$\frac{\partial H}{\partial G} = 0 \Rightarrow p(t) = n_G q(t) \tag{10}$$

$$\frac{dp(t)}{dt} = \rho p(t) - \frac{\partial H}{\partial k} \Rightarrow \frac{dp(t)}{dt} = [(\rho + \delta) - y_k]p(t) - n_k q(t) \tag{11}$$

$$\frac{dq(t)}{dt} = \rho q(t) - \frac{\partial H}{\partial \Omega} \Rightarrow \frac{dq(t)}{dt} = [(\rho + \zeta) - n_\Omega]q(t) - y_\Omega p(t) - (1 - \psi)\Lambda_\Omega \tag{12}$$

In the above expressions, U_c, n_G, y_k, n_k, n_Ω, y_Ω, and Λ_Ω correspond to partial derivatives of the various functions in the model with respect to the indicated variables. Transversality conditions apply:

$$\lim_{t\to\infty} p(t)e^{-\rho t}k(t) = 0 \tag{13}$$

$$\lim_{t\to\infty} q(t)e^{-\rho t}\Omega(t) = 0 \tag{14}$$

Because most of the derivatives in the optimality conditions correspond to relatively intricate nonlinear terms, it is not feasible to compute explicit expressions for the equations of motion of the problem's control variables. Thus, one is not able, as well, to deliver the steady state expressions of the endogenous variables. For numerical examples, including the one in the main text, one may determine the following set of steady state values: $(k^*, \Omega^*, c^*, G^*)$. The computation requires imposing conditions $\frac{dk(t)}{dt} = 0$, $\frac{d\Omega(t)}{dt} = 0$, $\frac{dp(t)}{dt} = 0$, $\frac{dq(t)}{dt} = 0$ to equations (5), (2), (11), and (12), and to use conditions (9) and (10) to suppress the shadow-prices from the obtained system of equations. Then, the system could be solved for the remaining four endogenous variables.