KUNGL. VETENSKAPS-AKADEMIEN
THE ROYAL SWEDISH ACADEMY OF SCIENCES

12 OCTOBER 2026



Scientific Background on the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2026

MICROECONOMIC FOUNDATIONS OF MACROECONOMIC STABILIZATION

The Committee for the Prize in Economic Sciences in Memory of Alfred Nobel

Microeconomic foundations of macroeconomic stabilization

1 Introduction

The stabilization of aggregate economic activity is a central goal of economic policy. Macroeconomic fluctuations are viewed to be undesirable to an important extent, especially since some downturns involve unemployment and a seeming failure of markets to fully use all available production factors; in addition, high and variable inflation hurts consumers and firms in their activities in a number of ways. Moreover, these fluctuations do not hit evenly across the population, as both labor market outcomes and prices tend to swing more for poorer households, thus making recessions times of heightened inequality. For these reasons, many regard carefully designed stabilization policy as one of the most important goals of economic policy overall.

How and when to stabilize is, however, challenging. Is it, for example, the right time now for the government to stimulate the economy to lower unemployment, or should it wait, and should they lower taxes or enact transfers, and for which consumers? Should the central bank increase the interest rate to curb inflationary pressures, and by how much? The questions, of course, presuppose that central banks and governments are, in fact, able to affect the way business cycles play out. But just how different interventions affect the macroeconomy is of great importance and is, in practice, key to the actual operation of economic policy.

The operation of macroeconomic stabilization policy since WWII appears to have worked rather well until the beginning of the 1970s, when it became apparent that the relationships between key aggregate variables, such as inflation, unemployment and GDP, were much less stable than previously thought. This undermined the paradigm at the time, since the conduct of stabilization policy had been based on the robustness of these relationships.

The instability cast serious doubts on the prevailing theories of stabilization—theories that were developed chiefly by John Maynard Keynes in response to the Great Depression almost half a century earlier. A new framework for macroeconomic stabilization policy needed to be built up, one that, like the Keynesian approach, also recognized the shortcomings of mar-

kets and, therefore, the need for government interventions. However, in contrast to the previous framework, the new one had to rest on an understanding of the fact that macroeconomic outcomes cannot simply be represented as stable relationships between aggregate variables. Rather, macroeconomic outcomes had to be seen as the cumulative result of actions taken by all the individual consumers and firms in the economy. The experience during the 1970s also convinced the profession that consumers and firms were far more forward-looking than had been previously thought: they pay careful attention to where both monetary and fiscal policy are heading and, hence, to expected inflation and future aggregate activity. Therefore, a new framework had to describe economic actors as purposefully reacting to government stabilization policy as well as to the fundamental factors causing the fluctuations in economic activity. Lastly, these fundamental factors should be expected to vary in nature over time: different recessions occur for different fundamental reasons, and may therefore also require different policy measures.

This year's prize recognizes two researchers—Robert E. Hall and Michael Woodford—for their key contributions toward a new macroeconomic framework for monetary and fiscal stabilization policy. In a seminal article from 1978, Hall showed that a theory of consumption smoothing—the idea that consumers prefer consumption to be stable over time—worked surprisingly well for understanding the movements of aggregate consumption in the data. The theory had elements in common with Keynesian consumption theory but was, at its roots, fundamentally different in that it relied on fully rational, forward-looking consumer behavior rather than mechanical rules of thumb. The theory in particular implied that consumption was not as easily manipulable by government policy as was previously thought, potentially thus limiting the effects of stabilization policy.

Through a series of subsequent papers, built on Hall's consumption framework as well as other important developments in macroeconomic theory during the 1970s and 1980s, Michael Woodford arose as the key architect behind the so-called New Keynesian model: a full macroeconomic stabilization framework that has become the main tool for policymakers today. In his 2003 monograph, Woodford summarizes as well as further develops the general setting. The original theory was chiefly aimed as a foundation for monetary policy analysis but subsequently also became the main tool for the analysis of fiscal policy. This model features departures from perfect competition, with firms charging markups over marginal cost, as well as sticky prices and wages preventing the efficient use of resources. Despite allowing consumers and firms in the model to adapt their behavior to changes in government policy, it shows how appropriately used fiscal and monetary policy can improve welfare by stabilizing the economy in the face of disturbances.

As the core model had been around for about a decade when the Great Recession suddenly

hit our economies in 2008, it turned out to be of great practical use, as it helped steer the economy onto the right path. This help came not only because the new framework helped us analyze the situation, but also by suggesting new policy instruments; we briefly discuss how the awarded research contributed to policymaking over the last quarter century in Section 5 below.

The specific research contributions underlying the prize were put together independently and with quite some time in between. The final macroeconomic stabilization framework also builds on contributions of many different researchers, including several Nobel Laureates in the Economic Sciences.

Hall's early research on consumption was motivated by its importance for economic fluctuations and stabilization policy.[1] A focus on consumption is natural: private consumption is by far the largest component of aggregate demand in the economy. This, and the fact that private consumption is determined not by the government or central banks, but by individual decisions made by millions of households, implies that a thorough understanding of the determinants of consumption is necessary for stabilization policy. For example, we first need to understand how consumption demand is affected by changes in disposable income or interest rates before we can analyze how consumption in turn affects aggregate economic activity, employment, and inflation.

A key complication in answering these questions is that we have strong theoretical reasons to believe that consumption decisions are forward-looking.[2] That is, the level of consumption depends not only on current income but also on the expectations people have about future income and other relevant variables, such as economic policy. For example, a change in disposable income (e.g., due to a tax cut or public subsidy) might be saved rather than consumed if it is expected to be temporary. But to the extent it is perceived as being long-lasting, it will likely lead to a much larger change in consumption and, hence, our macroeconomics aggregates.

Since these expectations cannot be directly observed in data, it appeared impossible to test the theory and quantify to what extent expectations drive consumption behavior. However, Hall (1978) provided a characterization of forward-looking consumption behavior that could be directly tested in data. The key insight is that the underlying theory implies that, given current consumption, future consumption should be impossible to predict. Very surprisingly to the profession, Hall then showed that this prediction was largely borne out in aggregate consumption data. This led to a "Eureka" moment of sorts for the macroeconomic community, since it removed a main obstacle for building rational, forward-looking behavior into models of economic policy. Hall (1978) was followed by Hall and Mishkin (1982), which addressed the individual consumption decisions of households underlying aggregate consumption. The au-

[1] This is for example stressed in Hall (1986).

[2] These ideas are consistent with the research awarded in past economic prizes, e.g., that by Milton Friedman and Franco Modigliani.

thors posited that many consumers were not able to smooth consumption over time, perhaps because they were resource-constrained and were unable to borrow, instead simply consuming their current disposable income: they live "hand-to-mouth". In particular, the authors could now estimate, using panel data on individual consumption, what fraction of aggregate consumption is forward-looking as opposed to hand-to-mouth. Hall's 1978 paper thus generated a new approach that became the foundation for modern research on the determinants of private consumption.[3]

Michael Woodford's early contributions, prior to his 2003 monograph *Interest and Prices: Foundations of a Theory of Monetary Policy*, consisted of a series of papers, all adding different key building blocks toward what later became the full-fledged New Keynesian model. Prior to Woodford's own papers and his book, influential insights by other researchers include Taylor (1979), Rotemberg (1982b), Calvo (1983), and Blanchard and Kiyotaki (1987). Among his own papers, Rotemberg and Woodford (1992), Woodford (1994), Woodford (1995), Rotemberg and Woodford (1996b), Rotemberg and Woodford (1996a), Woodford (1996), Rotemberg and Woodford (1997), Woodford (1998), Woodford (1999b), and Rotemberg and Woodford (1999) all constitute key steps forward; Julio Rotemberg passed away in 2017. Each paper thus adds an important component toward the final edifice, and further contributions by Woodford and others, many of direct practical relevance, followed in the early 2000s and later.

The New Keynesian model features price (and wage) setting with frictions, implying that prices are not changed every instant. Because of this, monetary policy can have significant real effects. At the same time, the model is built on microeconomic foundations, which makes it possible to avoid the so-called "Lucas critique": if a macroeconomic theory is built on reduced-form relationships that are chiefly correlations between macroeconomic aggregates in the data, then a policy change will most likely alter the individual behavior underlying the observed aggregate variables and hence change the nature of their correlations.[4] This would show up as instability of the macroeconomic relationships to the extent policy tried to systematically make use of them, something that had been all too clear during the 1970s. By modeling the underlying determinants of consumer and firm behavior, the new theory could thus be expected to be robust to policy changes; the Lucas critique and the historical background to the New Keynesian model is described at some length in Section 2 below.

Another important component in the core model is the absence of a monetary aggregate. This feature of the model was not merely a simplification or a premonition of a future "cashless society". More importantly, it made monetary policy analysis focus on *interest-rate rules*, which was a significant step away from the previous paradigm whereby central banks aimed to control

[3] Notable further contributions to the research on consumption include the work of Angus Deaton, the 2015 economics laureate.

[4] Robert E. Lucas, Jr. was awarded the economics prize in 1995.

money supply, and where the interest rate was merely a policy instrument toward this goal.

Through his work, Woodford thus provided a new foundation for interest rate-based policies where central banks affect the economy by setting targets for the nominal interest rate, an approach that is now adopted by the vast majority of central banks. Furthermore, he used this framework to analyze monetary policy from a welfare perspective, deriving optimal policy rules. Together with Rotemberg, Woodford also demonstrated how a New Keynesian model could be estimated based on economic data and, hence, used quantitatively. This was an important step, as it made it possible for policymakers to construct large econometric models for predicting and quantifying the effect of various policy-measures, greatly improving on the models based on the traditional Keynesian framework.

The simplest New Keynesian model has only three equations; nevertheless, it is a very powerful tool for understanding some key aspects of how monetary policy works. The first equation summarizes consumer behavior and is in fact Hall's theory of consumption. The second equation describes how profit-maximizing firms set prices when it is costly to change them, and the third and final equation captures the central bank's interest-rate policy. When analyzing monetary policy quantitatively using this framework, additional variables and equations are added; central banks often have medium- as well as large-scale versions tailored to their national economies in question. Similarly, finance ministries and government agencies analyzing fiscal policy build versions of the model with yet another, slightly different, focus.

With these models, fiscal policy and monetary policy can be jointly analyzed. For example, such models show that when a government aims to stimulate the economy using transfers targeted to households, the newly created consumption demand will generate upward pressures on prices, which in turn leads an independent central bank with a mandate to control inflation to raise interest rates and hence counteract the initial stimulus. The framework also highlights how different shocks should be met with different policy reactions, even though they may all move output in the same direction. Hence, the new framework also led to a new approach to policy analysis, whereby the deeper nature of shocks, and the details of how each shock is expected to propagate through the economy, came into focus. The framework has thus been adopted broadly and globally and is perceived to have been very helpful for policymakers also in generating better outcomes.

This document proceeds as follows. We first provide a background for both Hall's and Woodford's work in Section 2; this section explains how the new stabilization framework was fundamentally different from that built by Keynes but at the same time stood on the foundations of work from the early 1980s emphasizing the role of markets much more than that of government policy—the so-called Real Business Cycle (RBC) theory.[5] We then describe the laureates'

[5] Through this research, Finn Kydland and Edward C. Prescott were awarded the economics prize in 2004.

key contributions chronologically, first Hall's research on consumption and then the New Keynesian model. We briefly also take stock from the perspective of how the most recent recessions have been handled—in Section 5—and in Section 6 we conclude.

2 Macroeconomic modeling from Keynes and on

The Great Depression during the 1930s hit broadly and deeply across the economies of the western world. A particular case in point is the United States, whose economy had risen to become a globally leading growth engine, a feat that was believed to be due to its largely unregulated market-based economy. Now, however, the same market economy experienced a sudden and sharp collapse—apparently, a new framework was needed.

Out of the analysis of the fallout grew a particular theory, namely that proposed through the writings and broader influence of the economist John Maynard Keynes. Keynes's core notion was that markets had severe shortcomings and that these shortcomings could be made less severe, or even eliminated, by appropriate government interventions. In short, Keynes's view was that recessions were times when there was a *shortfall in the demand* for goods and services, both for private consumption and for investment. Thus, firms had the capacity to produce—there was no shortage of either workers or physical capital—but the buyers of their products held back, for reasons Keynes considered erratic. Thus, aggregate demand was seen as affected by the "animal spirits" of investors and consumer "confidence", notions that were left unexplained but nevertheless could generate economic downturns. A basic remedy, Keynes thought, was for the government to somehow make aggregate demand move up again.

It does not take a trained economist to raise a fundamental question here: why didn't firms lower their prices so as to attract the customers back? If this was because wage costs were high and lower prices would mean that firms made losses, why didn't wages adjust down, so as to make it possible to keep the workers employed? Keynes held that the basic reason was that prices and wages simply were rigid, i.e., there apparently were costs associated with changing them. Markets were thus not always capable of adjusting to demand shortages. Instead, a sudden demand shortage could be amplified as the laid-off workers experienced severe income losses and hence lowered their consumption demand further, leading to further lay-offs. Fortunately, we have not experienced another recession as deep as the Great Depression, when unemployment rose to very high levels and remained high from the early 1930s to much later in the decade. Still, and more generally, unemployment is a highly volatile variable, as it rises sharply in recessions, which is seen as one of the major challenges to policymakers. Keynes's theories received broad support and both fiscal and monetary authorities used their available instruments to steer demand back up when it appeared to be needed; too strong a stimulus was

undesirable because as the economy neared full capacity, further demand would simply raise prices and threaten price stability.

2.1 The standard IS-LM and AS-AD framework

Keynes's theory was exposited and became standard textbook material through contributions by John Hicks, one of the very first Laureates in Economic Sciences (in 1972). He captured the key macroeconomic relationships transparently with linear equations, summarized in the so-called IS-LM model; here, IS stands for Investment-Saving and represents the market for consumption and investment goods, whereas LM stands for Liquidity preference-Money supply and describes the money market. We describe a very simple version of it here.

Aggregate output is assumed to equal aggregate demand, given the presumption that there is unemployment and capital is not utilized fully. Aggregate demand, in turn, consists of consumption, investment, and government expenditures.[6] A simple Keynesian consumption function postulates consumption C to be an affine function of disposable income $Y - T$ (where T is net taxes), i.e., $C = \alpha + \gamma(Y - T)$, where the parameter $\gamma \in (0, 1)$ is the marginal propensity to consume. Given that supply is not constrained, it will simply equal what is demanded. Leaving investment, I, and government expenditures, G, unspecified, we can then equate gross domestic product, Y, to the sum of all demands: $Y = C + G + I$.

Combining this equation with the consumption function allows us to solve for the equilibrium output level as

$$Y = \frac{1}{1-\gamma}(\alpha + I + G - \gamma T). \tag{1}$$

Here $1/(1-\gamma) > 1$ captures the amplification alluded to above in the form of a multiplier: it describes how 1 billion US $ of government spending, G, translates into a higher output level. This multiplier crucially depends on the marginal propensity to consume, γ.[7] Together with a parametrization where $\gamma \approx 2/3$—representing the fact that roughly 2/3 of disposable income is consumed and the rest saved—we obtain a multiplier of about 3 from a G increase and a multiplier of 2 from sending out checks. Thus, in a deep recession, paying people to dig holes in the ground and fill them up again looks like a fiscal policy that actually works. To derive the IS relationship—that between output and interest rates—Hicks further assumed that investment depends negatively on the interest rate: equation (1) would then give output as a negative function of the interest rate. To complete the description of how output is determined in the

[6] We regard the economy as closed, i.e., imports and exports are zero.

[7] This multiplier does not only apply to government spending. Imagine a government sending all citizens a stimulus check (as was done in the US under the CARES Act in 2020): the check would correspond to a decrease in T and hence the corresponding multiplier is again proportional to $1/(1-\gamma)$, but now multiplied by γ itself, since the transfer check is not spent in full—only a fraction γ is. The closer this propensity is to one, the higher is the effect of a stimulus on output.

short run when prices and wages are rigid, it was also necessary to describe how interest rates—market rates that were not believed to feature any rigidity—were determined in the short run. For this, the LM curve was developed.

The LM relationship is simply a model of the "money market": households and firms are believed to have a demand for real money balances, i.e., the amount of goods that money can buy, M/P, where M is nominal money and P the price level, which is increasing in income Y and decreasing in the nominal interest rate, i: $M/P = L(Y, i)$.[8] That is, money is viewed as liquidity, and the alternative to holding liquid money, which does not bear interest, is to keep it invested at i. The LM curve is thus simply obtained when the demand for money equals supply: $L(Y, i) = \bar{M}/P$, where $\bar{M}$ is the nominal money supply, governed directly by the central bank. This equation defines a positive relation between Y and i, and together with the IS curve, the system gives us how Y and i are determined in the short run when prices are fixed, i.e., not yet adjusting.

Textbook stabilization policy based on the IS-LM model allows fiscal policy to affect demand, through changes in G or T, and it allows monetary policy to affect output through changes in the money stock. An increase in $\bar{M}$ will, since P is fixed in the short run, induce a lower i, for any given Y, and the lower interest rate will stimulate investment through the IS curve, leading to higher output in the economy.

The three-equation New Keynesian model, which will be described in detail below, replaces the IS curve with a similar relationship, but now based on a different theory of consumption—that proposed in Hall (1978). It replaces the LM curve with an interest-rate rule, i.e., a description of how the central bank conducts monetary policy by setting the interest rate directly, without mention of monetary aggregates.

Lastly, the New Keynesian model is a fully dynamic model, i.e., it describes how the economy's aggregates evolve over time, month by month. Therefore it includes a description of how prices are set: they are not just fixed, but move sluggishly—the third equation of the New Keynesian model describes this process.

The dynamic adjustment of prices is thus not a subject of analysis in the traditional IS-LM model; instead these are given by a separate, "medium-run", model, usually labeled the AS-AD (for Aggregate Supply-Aggregate Demand) framework. Here, the AS part is a traditional Phillips curve, describing how the rate of inflation is negatively related to unemployment: there is a tradeoff, whereby lower unemployment can be achieved at the cost of higher inflation. This relation between inflation and unemployment, first described by Phillips (1958), was also a cornerstone of the Keynesian apparatus.

The AD curve also relates quantities to prices and can be derived in different ways. One

[8] Here, thus, L is a function that is increasing in Y and decreasing in i. Take note that i is nominal but moves hand in hand with the real interest rate, since prices, and hence inflation, are fixed.

derivation is to use the LM curve but no longer consider P as fixed, and to replace i in it with i from the IS curve. The New Keynesian model's dynamic price adjustment equation thus replaces the traditional Phillips curve.[9]

2.2 Lucas, the break, and the RBC model

All the relationships in the IS-LM and AS-AD frameworks were viewed to be robust and possible to use for "counterfactual analysis", i.e., scenarios for policy or other aspects in the economy that could have been observed but had not been historically. This view had become conventional wisdom: the Keynesian paradigm was widely regarded as highly successful for several decades—up until the 1970s. At that time, large parts of the western world were caught in a stagflation period: high unemployment *and* high inflation. Thus, it looked like the Phillips curve had shifted outward, i.e., that any given level of unemployment was now associated with higher inflation than before.

Robert E. Lucas, Jr., the 1995 Laureate, then argued that the instability of the Phillips curve, and indeed of the expanded set of IS-LM and AS-AD equations that had become the empirical implementation of the Keynesian model, was unavoidable and a flaw of the entire paradigm. He specifically also criticized the Keynesian consumption function; his broader point—the Lucas critique—was that all the relationships of the Keynesian framework were merely vaguely motivated by theory, but not derived from deeper parameters that one could expect to be robust, i.e., not change, when policy changed. He in fact showed that the coefficients in the Keynesian equations *should* change as policy changed. He went on to suggest, along with leading economists at the time like Edward C. Prescott (2004 laureate), Thomas J. Sargent (2011 laureate) and Neil Wallace, that a wholly new framework was needed, one that built explicitly on microeconomic foundations: "deep parameters" had to be estimated and "structural relations" had to be identified at the micro level.

The idea then was to build a macroeconomic model describing the behavior of aggregate variables as an explicit aggregation of the actions of individual economic agents such as households and firms. These actions, moreover, cannot be expected to be well described as simple rules of thumb; instead the objectives of the agents, the environment they operate in, and the constraint they face must be modelled in a sufficiently realistic way, and then agents' decisions must be derived based on this information. As a consequence, the modeler would need to assume that consumers and firms make forward-looking decisions—on how much to work and how much to save. Then, if government policy were to be changed, the environment in which agents make their decisions would likely affect their expectations about the future.

[9] Interestingly, an impactful intermediate case was already postulated in the 1960s by the economics laureate of 2006 Edmund Phelps: the expectations-augmented Phillips curve. This relationship is closer to that in the New Keynesian model, though not formally derived based on how firms set prices in a dynamic environment.

The incorporation of such a mechanism would then potentially change the relation between aggregate variables in important ways. Lucas's most well-known example, from his 1972 paper (Lucas (1972)), was that the Phillips curve could shift, to the extent the central bank tried to trade off inflation against unemployment; this mechanism could explain how the attempts to lower unemployment during the 1970s by raising inflation merely raised inflation, without lowering unemployment. With microeconomic foundations, such changes could be predicted systematically instead of ignored, or assumed to be invariant to policy. At least in theory, then, the models would be immune to the Lucas critique.

The paper by Finn Kydland and Edward C. Prescott from 1982 precisely proposed a framework, the RBC framework, that they considered to be the proper path forward for macroeconomic theory; Kydland and Prescott shared the 2004 Prize. The RBC model was built explicitly on microeconomic foundations, with consumers and firms optimizing well-defined objectives and prices acting to clear markets under perfect competition. In short, they described a world in which markets do wonders and do not have the flaws that were central to the Keynesian paradigm. By the same token, there was no role for government in stabilizing the economy. Their theory, moreover, was quantitatively oriented by design: it addressed the macroeconomic data not just qualitatively, and it showed surprising success in accounting for business cycles. Of course, the RBC theory, due to its presupposition that markets work so well, generated a flood of criticism from the proponents of the Keynesian paradigm. The criticism did recognize the weaknesses of the IS-LM and AS-AD model, but simply thought that the RBC model was much too far from a realistic description of how our economies work. In particular, they focused on the RBC model failing to recognize features such as price and wage stickiness.

Notwithstanding the frustration with the RBC framework, some researchers did take up the challenge Lucas raised: they began attempts at constructing a new macroeconomic framework that was built on microeconomic foundations and was robust to the Lucas critique. Importantly, it also featured elements of market failure, however with explicit descriptions of these failures. A paper was published as early as 1982 on the topic, by Julio J. Rotemberg, followed by a number of similar, but also different, approaches.[10] Most people followed in Rotemberg's footsteps, and this year's Prize celebrates the accomplishments of the two researchers who provided crucial building blocks of a new framework for the analysis of stabilization policy: Hall in providing insights into the behavior of aggregate consumption, the key building block for the "New Keynesian IS curve", and Woodford for contributing many crucial components, not least the final construction, of the New Keynesian model.

[10] As an example of an alternative approach, in 1982 Peter Diamond (laureate of 2010) published a paper with a general-equilibrium search-theoretic approach to market failures (see Diamond (1982)).

3 Hall's study of intertemporal consumption behavior

This section begins with a brief description, in Section 3.1, of the state of the literature prior to Hall's 1978 paper. Section 3.2 then lays out the optimal consumption problem Hall focused on in his pathbreaking paper, and Section 3.3 describes the test he proposed there. The paper used aggregate data only, and Section 3.4 describes the follow-up research examining aggregate data. Section 3.5 then covers the research confronting microeconomic data, which began with Hall and Mishkin (1982), and Section 3.6 concludes.

3.1 Points of departure

In the 1970s, the macroeconomic literature on consumption was divided. First, parts of the literature followed Keynes (1936), who asserted: "The fundamental psychological law, upon which we are entitled to depend with great confidence both a priori from our knowledge of human nature and from the detailed facts of experience, is that men are disposed, as a rule and on the average, to increase their consumption as their income increases, but not by as much as the increase in their income." The classic Keynesian approach was then to simply postulate consumption to be a function of disposable income, e.g., an affine relationship

$$C = \alpha + \gamma(Y - T), \tag{2}$$

with parameters $\alpha > 0$ and $\gamma \in (0, 1)$. Selecting the parameters of such a function to fit the data typically delivers a relatively high marginal propensity to consume γ of around $2/3$, which in turn implies a significant multiplier of about 3 (see the discussion in Section 2).

This simple approach held some empirical promise: it implied consumption to be less volatile than output and, for the cross-section, a larger saving rate for richer households. However, it also faced some challenges, in particular matching the rather stable saving rate in the time series despite significant economic output growth (see Kuznets (1962); Simon Kuznets was a 1971 Laureate for the Prize in Economic Sciences). The most severe critique came from Fisher, Modigliani, and Friedman, who attacked the approach on theoretical grounds and insisted that consumption must obey intertemporal considerations, and from Lucas (1976), who explicitly took issue with specifying an ad hoc consumption function and questioned its usefulness for informing policy.

The main alternative approach to the Keynesian consumption function was based on Friedman's theory of "permanent income" (Friedman (1957); Milton Friedman was a 1976 Laureate for the Prize in Economic Sciences), which explicitly takes an intertemporal perspective under uncertainty and predicted consumption to be related in a stable way only to permanent

income, i.e., the expected annuity value of net worth or, formally, $C_t = \epsilon Y_t^P$, where ϵ is a constant. In contrast to the Keynesian perspective, the theory predicts the marginal propensity to consume out of transitory income changes to be small (in the order of magnitude of the interest rate of about 2%). Hence, whether one chooses this permanent-income route or the more traditional Keynesian approach to consumption will have first-order consequences for the analysis of stabilization policy.

The key challenge of the permanent-income perspective arose from the empirical implementation of the theory, since the concept of permanent income involves unobserved expectations about the future. The literature took the approach of using a distributed lag structure to "guesstimate" permanent income from a series of past income levels. An often adopted formulation was to use geometric lags (so-called Koyck lags). Muth (1960) had pointed out that such a geometric lag specification only captures the rational permanent-income forecast under a very specific income process, but his point did not seem to change the course of the literature: prior to the rational expectations revolution in the 1970s, research on consumption kept using ad hoc distributed-lag specifications without a theoretical basis, and therefore remained vulnerable to the Lucas critique.

3.2 A canonical intertemporal consumption problem

Hall (1978) marked the beginning of a new approach. Hall's paper is based on the following household problem: a household has capital income and stochastic labor earnings and decides each period how much to spend on consumption subject to budget constraints. Formally, the household problem reads

$$\max_{\{c_{t+\tau}, a_{t+\tau+1}\}_{\tau=0}^{T}} \mathbb{E}_t \left\{ \sum_{\tau=0}^{T} \beta^\tau u(c_{t+\tau}) \right\}, \tag{3}$$

subject to

$$a_{t+\tau+1} = (1+r)\left(a_{t+\tau} - c_{t+\tau} + w_{t+\tau}\right),\ \tau = 0, 1, \ldots T, \tag{4}$$

$$c_{t+\tau} \geq 0,\ \tau = 0, 1, \ldots T, \tag{5}$$

and a terminal condition

$$a_{t+T+1} \geq 0, \tag{6}$$

for a given initial wealth level a_t.[11] Here, we thus consider a household that plans for $T+1$ periods from period t onward and can save and borrow at a fixed interest rate r. Equation (3) states the objective function, i.e., "lifetime utility," as the geometrically discounted sum of

[11] The notation presumes that the sequences for consumption and asset holdings depend on the realized histories of earnings.

period utility levels $u(c_{t+\tau})$, where $\beta \in (0,1)$ is the discount factor. The period utility function $u(c)$ is assumed to be strictly increasing and strictly concave in its argument. The maximization is carried out subject to standard period budget constraints (4), which state that wealth next period, $a_{t+\tau+1}$, is equal to one plus the interest rate, r, times this period's wealth $a_{t+\tau}$ plus labor earnings $w_{t+\tau}$ minus consumption expenditures $c_{t+\tau}$. The interest rate $r > 0$ is assumed to be constant but labor earnings $w_{t+\tau}$ are stochastic and constitute the source of uncertainty in the problem, in line with the Fisher-Friedman-Modigliani literature.[12]

As earnings are stochastic, the expectations operator $\mathbb{E}_t$ shows up in the household objective: the household makes an optimal consumption decision in period t to maximize expected lifetime utility given all the information up until period t. Hence, the household is maximizing a von Neumann-Morgenstern expected utility function and $\mathbb{E}_t$ is the expected value operator given all the information revealed up to period t. For instance, by period t, a_t and w_t are known, but w_{t+1}, w_{t+2}, etc. are not, so the household has to build expectations over them to optimally choose today's consumption level. Finally, initial wealth is given and there is a terminal condition (6) on wealth being non-negative after the $T+1$ periods: the consumer cannot "die indebted".

The household problem (3)–(6) later became the canonical way of specifying intertemporal consumption decisions: the economics literature since the 1980s is building on some version of this formulation at its core. Hall was not first in specifying this canonical household problem but—as we will explain below—the implications he derived from this problem were groundbreaking and opened up for a completely new path toward testing the theory's prediction in data.

Often, macroeconomic models use an infinite time horizon, where a household is thought of as a "dynasty," in which case (6) is replaced by the so-called no Ponzi game condition.[13] Also, macroeconomic analysis—like the RBC model or the New Keynesian framework outlined below—usually adds a labor supply decision to the household problem and integrates such a household block with a production side that determines the equilibrium interest and wage rate endogenously; the above formulation, in contrast, is one in partial equilibrium and thus assumes the interest rate and wages to be exogenous.

Since the mid-1990s, macroeconomic models often also consider formulations with borrowing constraints of the sort $a_{t+\tau} \geq \bar{a}$, $\forall \tau$; the value $\bar{a}$ can either be derived as a natural borrowing limit, i.e., the maximal borrowing level such that the consumer can pay back the loan regardless of the realization of shocks, or can be set to an exogenous value. In such models, as in the

[12] Macroeconomics often allows the interest rate to be time-dependent and stochastic as well—more on this below, when we discuss Hall (1988). As a simplification, Hall (1978) abstracts from such fluctuations in the interest rates, in part surely motivated by the fact that the period he studied was characterized by a relatively stable real interest rate.

[13] This condition, $\lim_{T\to\infty} \left\{(1+r)^{-(T+1)} a_{t+T+1}\right\} \geq 0$, rules out pyramid-like borrowing schemes and is often formalized to hold "almost surely".

problem above, labor earnings constitute the main source of uncertainty the consumer faces. In sum, even though the present macroeconomic literature may add additional features to the analysis, the household problem (3)–(6) is the backbone and is still nested in more modern formulations.

3.2.1 Optimal consumption behavior

As we assume the period utility function to be strictly increasing in the argument, there will never be satiation and condition (6) will hold with strict equality. Furthermore, let us assume that the solution is interior, i.e., that the non-negativity constraints (5) do not bind; this can be shown to be true if one imposes so-called Inada conditions on the utility function. We can then consolidate the period budget constraints (4) and restate the household problem more compactly as

$$\begin{aligned} &\max_{\{c_{t+\tau}\}_{\tau=0}^{T}} \mathbb{E}_t \left\{ \sum_{\tau=0}^{T} \beta^\tau u(c_{t+\tau}) \right\}, \\ &\text{subject to} \quad \sum_{\tau=0}^{T} \frac{c_{t+\tau} - w_{t+\tau}}{(1+r)^\tau} = a_t, \end{aligned} \tag{7}$$

for all sequences of shocks, and a given initial wealth level a_t. In this problem, the (consolidated) lifetime budget constraint simply states that the discounted sum of consumption expenditures has to be equal to the discounted sum of lifetime labor earnings (plus initial wealth). Problem (7) is precisely the version of the household problem Hall (1978) specifies, analyzes, and, most importantly, brings to the data for testing under the assumption that utility is quadratic in consumption.

Solving the household problem (7) results in the following first-order conditions:

$$u'(c_t) = (1+r)^\tau \beta^\tau \, \mathbb{E}_t \left\{ u'(c_{t+\tau}) \right\}, \forall \tau = 1, 2, \ldots, T. \tag{8}$$

Equation (8) is the so-called stochastic consumption Euler equation. The Euler equation is an intertemporal optimality condition trading off consumption across different time periods. By forgoing a marginal unit of consumption in period t the household incurs a utility decrease of $u'(c_t)$, whereas it allows the household to obtain additional $1+r$ consumption units in the next periods yielding additional $(1+r)\beta \, \mathbb{E}_t \left\{ u'(c_{t+1}) \right\}$ in terms of expected utility from consuming tomorrow; for τs above 1, the trade-off involves saving and discounting over more periods, and hence both $1+r$ and β are raised to the power τ. Optimal behavior is attained when such a marginal change in savings gives zero net utility gains at all points in time, while also satisfying the budget constraints.

As there is diminishing marginal utility, the Euler equation reveals a motive for consumption smoothing. The Euler equation implies that the household optimally plans the marginal utility of consumption to change at a constant (gross) rate $((1+r)\beta)^{-1}$, and that the household as a consequence does not plan for any swings in the consumption level around this trend. To illustrate this further, let us assume that $\beta(1+r) = 1$: then the household plans to first order to stabilize the consumption level across all periods, irrespective of how volatile labor earnings are or of whether or not earnings tend to generally grow over the life cycle.

This is not to say that consumption will not fluctuate ex post, given that shocks to the income process do materialize. However, the key prediction is that the household does not plan for such consumption fluctuations ex ante.[14] This prediction is in line with the permanent-income hypothesis by Friedman (1957), though here it is formalized based on rational expectations. As we will see below, this formalization makes permanent income explicit. The theory makes a sharp distinction between the consumption response to anticipated and unanticipated income changes. A worker who is informed that they will be laid off in three months should cut back on consumption immediately and consumption should not further change once the contract indeed expires and income drops in an anticipated way.

The stochastic Euler equation (8) characterizes the change(s) in consumption between any two periods but does not pin down the consumption level. Nevertheless, the Euler equation makes a very sharp prediction by relating today's marginal utility of consumption to the interest rate, the rate of time preference and the expected future marginal utility of consumption. In particular, and remarkably, the level of initial wealth and the precise path of earnings between two periods should not affect the expected change in marginal utility of consumption between two periods. This orthogonality condition is exactly what Hall exploited in his famous test.

To retrieve the consumption level, the lifetime budget constraint is required in addition to the Euler equations. To illustrate how this works, let us consider the case where $\beta(1+r) = 1$ and where the utility function is quadratic, i.e., $u(c_{t+\tau}) = -\frac{1}{2}(\phi - c_{t+\tau})^2$.[15] These assumptions allow the derivation of a closed-form solution for the consumption level as follows:

$$c_t = \frac{r}{1+r-(1+r)^{-T}} \left(a_t + w_t + \mathbb{E}_t \left\{ \sum_{\tau=1}^{T} \frac{w_{t+\tau}}{(1+r)^{\tau}} \right\} \right). \tag{9}$$

This consumption function is thus based on the permanent-income hypothesis under rational expectations. The propensity to consume out of wealth or permanent (expected) lifetime earnings, then, just depends on the interest rate and the time horizon T. As the time horizon

[14] If the consumer had access to complete insurance against earnings shocks, also the ex-post swings in consumption could and would be eliminated.

[15] As the quadratic utility function neither features an Inada condition nor globally positive marginal utility, one has to place further restrictions on the earnings process (and initial wealth) such that the "bliss" consumption level, ϕ, is never reached and such that the non-negativity conditions (5) do not bind.

goes to infinity, this propensity becomes simply $\frac{r}{1+r}$, which implies—for reasonable values of the real interest rate of say 2% per year—a very low value. As discussed in Section 2 it makes a crucial difference for the effectiveness of stabilization policy whether the marginal propensity to consume is 1/50 or 2/3.

One might be tempted to try to construct an empirical test based on this consumption function (9), but this is challenging: unlike the traditional Keynesian consumption function (2), the rational-expectations, permanent-income consumption function, as represented by (9) in the special case of quadratic utility, is based on forward-looking behavior. Equation (9) in particular illustrates that consumption at t is a function of the period-t expectation of the present value of future lifetime labor earnings, and this expectation is not observed by the econometrician.[16] Even worse, it seems impossible to reject the theory, as one could always construct expectations about future labor earnings that would rationalize any observed consumption level c_t today based on the rational-expectations, permanent-income consumption hypothesis. As a consequence, the challenge of deriving a test of the theory from (9) seemed unsurmountable.

3.3 The test

Hall's idea was to construct a test directly based on the stochastic consumption Euler equation. He thus rewrote the stochastic consumption Euler equation (8) as follows:

$$u'(c_{t+1}) = ((1+r)\beta)^{-1} u'(c_t) + \epsilon_{t+1}, \forall t, \tag{10}$$

where ϵ_{t+1} is a disturbance term with $\mathbb{E}_t\{\epsilon_{t+1}\} = 0$. This is a significant step towards constructing a test, as (10) constitutes an equation with an error term that can potentially be brought to the data and tested in a systematic way. Hall's derivation of (10) is based on the insight that the expectation of a variable can be replaced by its realization and an expectational error, which can be folded into the error term.[17]

However, (10) is still cast in terms of marginal utility, which is not observed. In order to derive a sharp test in terms of (observed) consumption that could be implemented straightforwardly with a linear regression, Hall further assumed a quadratic utility function: $u(c_{t+\tau}) = -\frac{1}{2}(\phi - c_{t+\tau})^2$. The case of quadratic utility is of particular interest, as the resulting—linear—

[16] Of course, one could replace "permanent income" with an ad-hoc function of past income levels, e.g., using a distributed-lag specification, but such a specification would not be theory-consistent and therefore does not constitute a rigorous or convincing test of the theory.

[17] The same expectational-error substitution is already suggested in McCallum (1976a) and McCallum (1976b). This property later led to a burst in the empirical literature which exploited it for estimation purposes. The methodological development culminated in the generalized method of moments (GMM), which originated in Hansen (1982) and was awarded the Prize in Economic Sciences in 2013. The expectational-error substitution to recover structural Euler equation parameters has, however only been exploited after Hall (1978); see the early study by Kennan (1979) and the series of early applications to the consumption Euler equation in Mankiw (1981), Hall (1988) (which existed as a working paper as of 1981), Hansen and Singleton (1982), Hansen and Singleton (1983), MaCurdy (1983), and Shapiro (1984).

Euler equation can be seen as a first-order approximation of any continuously differentiable period utility function for small enough shocks. Hall's application of the test uses aggregate consumption data, which only show minor movements from year to year, making this approximation argument relevant.

With quadratic utility, (10) implies a linear relationship

$$c_{t+1} = \phi\left(1 - ((1+r)\beta)^{-1}\right) + ((1+r)\beta)^{-1} \cdot c_t - \epsilon_{t+1}. \tag{11}$$

The permanent-income hypothesis suggests one could estimate a conditional expectation $\mathbb{E}\left\{c_{t+1}|c_t, x_t\right\}$, where x_t is a vector of data available in period t, and then test the hypothesis that the conditional expectation is actually not a function of x_t. According to the theory, *none* of the controls that is in the information set of period t should—after controlling for the current level consumption and an intercept—have any predictive power for future consumption. After Hall (1978), the literature refers to his test as the "random walk consumption hypothesis test." This term comes from the simple case, where we set $(1+r)\beta = 1$, and the theory simply predicts

$$\mathbb{E}_t\{c_{t+1}\} = c_t, \tag{12}$$

i.e., it predicts consumption to follow a so-called martingale process (which is a generalized form of a random walk).

A strong assumption underlying the consumption framework above is that the household is assumed to be able to borrow and lend without limit and at a constant interest rate. Hall (1978) proposed and briefly discussed an alternative extended specification called rule of thumb, where a share μ of aggregate disposable income accrues to households which are liquidity-constrained in an extreme manner: they consume "hand to mouth." That is, this group of households would like to consume more but cannot borrow, and their consumption therefore follows contemporaneous disposable income one for one.[18] In this case, under quadratic utility and $(1+r)\beta = 1$, the change in per capita consumption becomes

$$c_{t+1} - c_t = \mu(y_{t+1} - y_t) - (1-\mu)\epsilon_{t+1}, \tag{13}$$

where y is disposable income per capita and ϵ_{t+1} is again a mean zero disturbance. Hence, with some liquidity-constrained households, i.e., $\mu > 0$, the test would show "excess sensitivity" of consumption with respect to temporary income changes, whenever income y fluctuates and is serially correlated. Similarly, under any traditional Keynesian consumption function that assumes a tight connection between contemporaneous disposable income and consumption,

[18] Another way to see this is as a special case of the Keynesian consumption function with $\alpha = 0$ and $\gamma = 1$.

the hypothesis (12) will be rejected. Hall also showed how the ad hoc empirical implementations of Friedman's permanent-income hypothesis that were predominant at the time—and that proxied for permanent income by flexible forms of distributed lags of income (without deriving permanent income under rational expectations)—imply predictability of consumption changes and a rejection of his test. Hence, it is really only the rational-expectations version of the permanent-income theory that leads to no rejection of the test.

The null hypothesis of Hall's test ran strongly counter to the predominant views and essentially all the empirical work on consumption at the time. Whereas a good empirical fit of model-derived equations was often taken as a validation of a theory, here data that are impossible to predict presented a metric of success of the rational-expectations, permanent-income hypothesis.

Strikingly, Hall indeed provided evidence that after controlling for today's consumption, disposable income and arbitrary lags thereof only have minimal predictive power for future nondurable consumption in aggregate US data at a quarterly frequency. The finding amounted to key evidence in favor of optimizing, forward-looking behavior.

Hall's findings came as a major surprise to the profession, not only because of the method for testing, but also because the results were found in aggregate data: a theory that actually applied on the level of individuals, and was hard to test even for individuals, actually seemed to work after aggregating across all the millions of different households in the economy. Soon after, there was widespread agreement that Hall's study amounted to key evidence in favor of the theory, which made it a cornerstone theory for understanding consumption.

In retrospect, the demonstration that the rational-expectations version of the permanent-income hypothesis suggested a strong form of unpredictability of (changes in) consumption also marked a methodological turning point for macroeconometrics and applied econometrics more broadly. There is a parallel to the analysis of asset prices here: the efficient-markets hypothesis suggests a strong form of unpredictability of price changes—see Fama (1970); Eugene Fama received the 2013 Prize in Economic Sciences. However, in financial markets, some agents with rational expectations may arbitrage away any predictability in stock market prices, so this kind of unpredictability is less surprising than in the case of consumption, where no such arbitrageur is present. It was therefore a real eye-opener that Hall showed how difficult it actually was to reject the random walk hypothesis in aggregate consumption data.

Hall's work also led rational expectations-based consumption theory to become a cornerstone of macroeconomic modeling. First, it was incorporated as a key element in the RBC literature, complemented with endogenous labor supply, implemented by adding hours worked, or leisure, in the utility function u. Second, the New Keynesian IS curve is in fact just a log-linearized version of the first-order condition discussed above (with a time-varying real interest

rate) and it is also a crucial component of this paradigm. Third, since the 1990s, macroeconomics has increasingly shifted toward modeling aggregate consumption as derived from many heterogeneous households. Hall's insights, and especially those that use individual consumption data, which we will discuss below, turned out to be crucial also for these recent developments.

It should be noted that, even in Hall's first tests, the non-rejection of the random walk hypothesis was not without some qualifications. Hall (1978) documented that the asset market valuation has some (though moderate) predictive power for changes in consumption, which is inconsistent with the pure form of the theory. This is also the general finding of the large follow-up literature on which we comment below. In particular, individual household-level data show a clear pattern of excess sensitivity of consumption with respect to temporary income changes. Pioneering in this respect is Hall and Mishkin (1982), which we discuss below, where an extended theory that allows for hand-to-mouth individuals is estimated in the household level data. The general pattern of "excess sensitivity" and high marginal propensity to consume of some households in turn was key for the literature to guide how the theory should be modified, e.g., by adding borrowing constraints to the canonical household problem (3)–(6).

3.4 Follow-up literature, using aggregate data

We first describe some immediate applications and tests that followed Hall's approach. We then describe how the literature evolved thereafter.

3.4.1 Follow-up tests of the rational-expectations, permanent-income hypothesis

An emerging literature (re-)ran versions of Hall's tests using data from the US and elsewhere; see, e.g., early international evidence in Johnson (1983) for Australia, Cuddington (1982) for Canada, or Daly and Hadjimatheou (1981) and Muellbauer (1983) for the UK. These studies typically identified some variables for which the hypothesis is rejected at standard levels of statistical significance (similar to asset prices in Hall (1978)). In an important early example using aggregate US data, Flavin (1981) documented evidence of excess sensitivity of consumption to income when including more lags of income into the test.[19] Hall's analysis was also extended in important dimensions, e.g., by allowing for a changing real interest rate, in Mankiw (1981), durable consumption goods, in Mankiw (1982), or in both, as in Mankiw (1985a).

The derivation and implementation of the test in Hall (1978) are based on a quadratic utility function, or at least the presumption that a Taylor expansion would be a good approximation

[19] As will be explained below, the analysis of Flavin (1981) went far beyond a simple extension of Hall (1978) by additionally explicitly modelling and estimating income processes. Additionally, Flavin (1981) also clarified the discrepancy between the findings by Hall (1978) and Sargent (1978), who provided early evidence rejecting the rational-expectations, permanent-income hypothesis.

of the nonlinear consumption Euler equation. In this case, when non-negativity constraints do not bind, the consumption function (9) is obtained in closed form. This setting imposes certainty equivalence, i.e., whether labor earnings risk is large or small does not affect consumption behavior (see the consumption function (9)).

Zeldes (1989b) added a word of caution on such rejections of the hypothesis based on certainty equivalence. If preferences are not quadratic, the martingale property still holds for marginal utility, but this does not exactly translate into a martingale for the level of consumption. Zeldes (1989b) thus studied a constant relative risk aversion (CRRA) utility function, computed the resulting consumption function numerically, and showed that the deviation from certainty equivalence may explain the finding of excess sensitivity in the literature.

3.4.2 Adding structure to the income process

The literature proceeded quickly beyond simply re-running Hall's test. A strength of Hall's test is that the information set of the agents and the precise income process do not have to be specified. However, this approach comes with the limitation that with such a test, if rejected, no statement about the quantitative significance of the departure, nor of the size of the marginal propensity to consume, can be made. For this reason, the following literature took further steps to move beyond testing by adding assumptions about the income process in order to back out structural parameters. Flavin (1981) was pioneering in this respect, where she showed how permanent income and an explicit structural consumption function can be calculated under a specific autoregressive-moving average (ARMA) process for income. Flavin (1981) defined the response of consumption to current income, beyond that attributable to the role of current income in signaling changes in permanent income, as excess sensitivity of consumption to current income and estimates this to be of the order of magnitude of 0.36. Furthermore, she estimated the marginal propensity to consume out of current income to be 16.75%.

Many papers followed the approach presented by Flavin (1981) and explicitly modeled and estimated the income process, thus studying consumption and income jointly by taking the cross-equation restriction between the two processes seriously. This led to influential studies of the variability in aggregate consumption and income documenting systematic deviations from the rational-expectations, permanent-income hypothesis benchmark in Hall (1978), e.g., the "excess smoothness" of consumption in Campbell and Deaton (1989): when there appear to be large changes in permanent income, consumption should respond, according to the theory, and this line of work argues that it does so too little in the data to be consistent with the theory.

Campbell and Mankiw (1989) build on the deviation from the benchmark briefly discussed by Hall (1978), where a fraction μ follows a rule of thumb as opposed to optimizing behavior. Then the change in consumption is given by (13). They use this equation to estimate the share

of hand-to-mouth households μ in aggregate data and test whether it is significantly different from zero. Recognizing that the error term ϵ_{t+1} is likely to be correlated with $y_{t+1} - y_t$, they adopted an instrumental variable approach.[20] In a baseline, Campbell and Mankiw (1991) use lagged changes in consumption as instruments and estimated the fraction of hand-to-mouth consumers to be around 0.5 and statistically different from zero in aggregate data from Canada, France, Japan, Sweden, the UK and the US.

Generally, the same orthogonality condition that is inherent to the stochastic Euler equation and underlies the test in Hall (1978) can also be exploited for estimation purposes. The key idea is to recognize that, under rational expectations, the intertemporal first-order condition will hold in expectation, even conditional on past variables. This means that parameters arising in these first-order conditions can be estimated by instrumental variables using past values as appropriate instruments, without a complete specification of an economic model, i.e., without specifying the precise source of uncertainty. After the breakthrough research in Hall (1978), the literature shifted away from further attempts to directly estimate reduced form "consumption functions"—an exercise vulnerable to the critique by Lucas (1976)—and instead turned to estimating "deep" (structural) parameters based on the intertemporal optimality condition exploiting the orthogonality condition described above.

This literature also benefited from the generalized method of moments (GMM) methodology developed by Hansen (1982), which allowed for more general utility functions than quadratic utility and a non-constant interest rate. The influential work by Hansen and Singleton (1982) and Hansen and Singleton (1983) provided estimates of the preference parameters based on this approach.

Another early example of the estimation of preference parameters is Hall (1988) (a study that had existed as a working paper since 1981). This paper extended the analysis in Hall (1978) to a changing interest rate and studied its implications for consumption. Hall considered a constant relative risk aversion (CRRA) period utility function, $u(c) = \frac{c^{1-\sigma}-1}{1-\sigma}$, and used the Euler equation to estimate the parameter σ. Hall also clarified that the $1/\sigma$ object that he estimated captures the intertemporal elasticity of substitution. Hall (1988) documented a weak relationship between the (expected) interest rate and consumption growth in the data, resulting in a very low estimate of the intertemporal elasticity of substitution of consumption of about 0.1, thus representing a strong smoothing motive for consumers. This finding is confirmed in, e.g., Campbell and Mankiw (1989).

Hansen and Singleton (1983), who used the same preference specification but focused on estimating the coefficient of relative risk aversion instead, found a small value for σ implying less curvature of the utility function. The discrepancy between the findings has to do with weak

[20] Again exploiting the orthogonality condition inherent to rational-expectations models, any variable that is in the information set at a particular point in time is potentially a valid instrument for future income changes.

instruments, and Attanasio and Weber (1993) challenged both estimates on the basis of an aggregation bias. Yogo (2004) investigated the issue further and confirmed the finding by Hall (1988) of a very low intertemporal elasticity of substitution when correcting for the weak instrument problem.

3.5 Microeconomic data

Hall (1978) and the researchers behind the initial follow-up literature focused their analyses on aggregate data, which were indeed the only data readily available to researchers at the time. Aggregate data have several disadvantages, however. The time series variation of aggregate income that can be exploited empirically is severely limited and the sample size becomes relatively small, restricting the statistical power of the analysis. Moreover, the theory models the consumption decision of a household, and the theory's predictions should therefore be tested at the micro household level rather than the aggregate level. The stochastic Euler equation that holds at the individual level only translates to the aggregate level under additional assumptions on the functional form of preferences. If these assumptions are not met or there is a systematic change in the household composition of the population, estimates that are based on aggregate time-series data run the risk of an aggregation bias. As a consequence, the modern empirical literature on consumption almost exclusively analyzes household-level panel data. A pioneering paper in this literature was Hall and Mishkin (1982).

3.5.1 Hall and Mishkin (1982)

Hall and Mishkin (1982) were the first to use microeconomic household level consumption data to test and study the rational-expectations, permanent-income hypothesis. As previously studied by Flavin (1981), Hall and Mishkin specified an income process and studied the movement of (food) consumption and income jointly, but now at the micro household level as opposed to on the aggregate level. More specifically, they used data from Michigan's Panel Study of Income Dynamics (PSID) on food consumption expenditure of 2,309 families between 1969 and 1975.[21]

Hall and Mishkin (1982) first documented a higher marginal propensity to consume out of permanent income than out of transitory income movements, in line with the rational-expectations, permanent-income hypothesis. However, when they ran the test in Hall (1978) at the microeconomic household level, they rejected the permanent-income hypothesis and documented an excess sensitivity of consumption to current income. The marginal propensity to consume out of transitory income shocks was also found to be of an order of magnitude higher than implied

[21] Earlier studies had used the PSID data to estimate the income elasticity of food consumption, e.g., Hymans and Shapiro (1976), but prior to Hall and Mishkin (1982) no study had confronted the rational-expectations, permanent-income hypothesis with the panel data.

by the benchmark rational-expectation, permanent-income theory. Finally, Hall and Mishkin augmented their theory (as briefly suggested in Hall (1978)) to allow a fraction of consumption to move hand in hand with current income and estimate this hand-to-mouth fraction in the panel data to be 20% (whereas the other 80% follow the prediction of the rational-expectations, permanent-income hypothesis).

The methods Hall and Mishkin applied, such as the maximum likelihood estimation of a whole system of equations, set a standard for the analysis of household consumption data that persists until today. Hall and Mishkin were also careful about numerous difficult details arising from consistently connecting the theory to the household-level micro data.[22] A large body of literature followed in their footsteps and estimated jointly income and consumption processes in household panel data—see Shapiro (1984) and Zeldes (1989a) as early examples—to inform the macroeconomic modeling of consumption behavior.

3.5.2 Follow-up estimates and tests from the applied microeconomic literature

Since the pioneering study by Hall and Mishkin (1982), the measurement of household-level consumer expenditure has expanded significantly. Large-scale consumer expenditure surveys, consumption imputed from income and wealth records, scanner/barcode panels, or financial transaction records allow the more recent empirical literature to go beyond the study of food expenditure, expanding further on the work by Hall and Mishkin (1982).[23] The result is a large body of work based on microeconomic data that significantly advanced our understanding of consumption behavior and that is used to discipline macroeconomic models of business cycles.

The rational-expectations, permanent-income hypothesis makes the very sharp predictions that households should adjust their consumption in response to unanticipated income shocks, but not to anticipated ones. A modern design-based empirical literature therefore sets out to test these predictions by exploiting quasi-experimental evidence, often coming from payments in the tax and transfer system or anticipated seasonal fluctuations in earnings.[24] This literature confirms the key findings of Hall and Mishkin (1982).

First, there is clear evidence of excess sensitivity of household level consumption to income changes, and the average estimates of the marginal propensity to consume out of transitory income changes are significantly higher than suggested by the rational-expectations, permanent-income hypothesis. Johnson et al. (2006) estimated an average marginal propensity to consume

[22] Some of the issues involve the treatment of household heterogeneity, measurement error in food expenditure, time aggregation due to the timing of the survey questions, etc.

[23] The PSID also added more and more consumption categories in 1999 and 2005, going beyond the initial coverage of food and housing expenditures.

[24] As examples of the former, Parker (1999) studied the effects of predictable changes in the social security tax, Souleles (1999) exploits predetermined income tax refunds, whereas Hsieh (2003) and Kueng (2018) exploited large anticipated payments from the Alaska Permanent Fund. As examples of the latter, Paxson (1993) and Browning and Collado (2001) based their study on large seasonal fluctuations in labor income.

out of the US 2001 rebates to be about 20–40% over three months, whereas Parker et al. (2013) documented the marginal propensities to consume out of the 2008 US stimulus payments to be 12–30%. Boehm et al. (2025) ran an experiment in France and documented an average marginal propensity to consume of 23% over a month. The meta-analysis by Havranek and Sokolova (2020) summarized the literature and found an average marginal propensity to consume of 21% among micro-estimates in the literature that exploit plausibly exogenous variation. After correcting for publication bias, Havranek and Sokolova (2020) reported a mean marginal propensity to consume of 11%.[25]

Second, the literature has robustly documented large heterogeneity in the marginal propensity to consume across consumers with some households showing close to hand-to-mouth type behavior. The estimation by Hall and Mishkin (1982) of the rule-of-thumb share can be seen as a clear forerunner of this finding.

Third, the literature also documents clear evidence in line with the rational-expectations, permanent-income hypothesis in particular when studying large salient payments as opposed to smaller, irregular ones. For example, Paxson (1993), Browning and Collado (2001), and Hsieh (2003) found considerable support of the permanent-income hypothesis. Even though Kueng (2018), who also studied payments from the Alaska Permanent Fund as did Hsieh (2003), showed considerably more mixed evidence, the implied welfare losses due to the deviations from the permanent-income hypothesis are very small. Households clearly smooth consumption to transitory income shocks (see, e.g., Blundell et al. (2008)) and consumption clearly reacts to news (see, e.g., Agarwal and Qian (2014) and Hendren (2017)). In sum, even at the microeconomic household level, changes in significant components of consumption are remarkably hard to predict.

3.6 Concluding comments

Hall (1978) created and tested a benchmark of the rational-expectations, permanent-income model, which had long remained a mere conjecture. After Hall's work, the theory has been cemented in macroeconomics.

Hall's contributions were important for the literature in many ways. First, Hall showed a path forward for systematic empirical testing of the rational-expectations, permanent-income hypothesis. This showed the usefulness of the stochastic Euler equation for estimation purposes. Second, even though refined versions of Hall's test typically show systematic deviations from his benchmark, particularly in microeconomic data, a significant component of consumption

[25] Havranek and Sokolova (2020) also reported systematic heterogeneity in the estimated marginal propensities to consume depending on the definition of consumption: broad measures of consumption that include durable goods tended to result in higher estimates, whereas narrower measures like non-durable consumption or food tended to lead to lower ones.

is nevertheless surprisingly hard to predict. This result suggested that the rational-expectations, permanent-income hypothesis had some clear empirical promise and provided a framework on which to build further, and the ensuing macroeconomic literature used it as a cornerstone. Third, Hall's simple benchmark enabled the systematic documentation of deviations, such as the excess smoothness of consumption or its excess sensitivity to income shocks, and the estimation of the fraction of hand-to-mouth households, as was carried out by Hall and Mishkin (1982). This led to important developments and modification of the theory, e.g., the inclusion of borrowing constraints. Fourth, Hall and Mishkin (1982) spurred many studies of the comovement of income and consumption in microeconomic data, which has greatly enhanced our understanding of the distribution of marginal propensities to consume and the estimation of policy-relevant preference parameters with direct consequences for stabilization policy. Such microeconomic estimates also form the empirical basis for the large literature on models with heterogeneous consumers.

The rule of thumb concept introduced in Hall (1978) was a precursor of a theoretical literature that added borrowing constraints to the macroeconomic analysis. This literature, which began with Huggett (1993) and Aiyagari (1994), can be seen as extensions of the RBC and New Keynesian frameworks to jointly address macro- and microeconomic data and helped bridge different sub-disciplines in macroeconomics.

4 Woodford and the New Keynesian model

We now describe the New Keynesian (NK) model. Its construction, a major intellectual achievement, has many dimensions, so it is impossible to describe it in full here.[26] Therefore, we focus on the "three-equation model," which is the framework in its very simplest form. It was, in fact, quite a surprise to the profession that the simple traditional Keynesian IS-LM/AS-AD framework could be replaced by a similarly compact new model that rested on full microeconomic foundations. This simplicity has made the model broadly accessible and very valuable in policymaking at central banks around the world.

Here we describe the New Keynesian model, its inner workings, and its applications in a real-world setting. Section 4.1 lays out the three key equations and discusses their interpretation and relation to the traditional Keynesian model. Section 4.2 provides the most important details underlying the model: its microeconomic foundations. Section 4.3 illustrates the workings of the model in simple examples and Section 4.4 discusses key further elements behind the development of the model. Section 4.5 discusses extensions, gives an example of an application to Sweden, and highlights some weaknesses of the framework.

[26] For a self-contained yet comparatively compact treatment, see Azzimonti et al. (2026).

Before describing the model in detail, let us briefly point to important early papers that were central for its development. The 1976 Lucas critique was not immediately embraced by all, but many economists who perceived Keynes's broad insights as very valuable did propose specific microfoundations that could resurrect demand management. First, as for papers modeling price and wage stickiness, we note Gray (1976), Phelps and Taylor (1977), Fischer (1977), Taylor (1979), and Taylor (1980). Note that these papers were all written before the RBC framework appeared and thus were not a reaction to it; their underlying real macroeconomic frameworks are different. These studies, moreover, assumed fixed rules for wage rigidities. Calvo (1983) then turned the attention to price stickiness; Mankiw (1985b) later allowed for menu costs, i.e., choice, and argued that small such costs could have large aggregate effects. Akerlof and Yellen (1985) made similar conceptual points.

The first paper that introduced price stickiness into fully dynamic models of price setting—the core of the New Keynesian model—is Rotemberg (1982a), which thus appeared simultaneously with the RBC paper by Kydland and Prescott (1982). The setting Rotemberg entertained has monopoly power and quadratic costs of changing prices; he considered monetary injections and showed they had real effects. Rotemberg (1982b) addressed the data and estimated a model with price stickiness, finding support for it. A number of other models where monetary injections had real effects were then introduced, including by Rotemberg (1987). Svensson (1986) and Blanchard and Kiyotaki (1987) studied models of monopolistic competition and showed how it, together with price rigidities, can allow monetary aggregates to raise aggregate output.[27]

Rotemberg and Woodford (1991, 1992) then argued that nominal rigidities are likely key (sufficient, though not necessary) for models where output is demand-driven to match the data. They argued that, under plausible conditions, one needed a countercyclical price-over-marginal-cost markup, or more generally a countercyclical wedge between firms' marginal product of labor and households' marginal substitution between consumption and leisure. Intuitively, with slow-moving prices, as output rises, prices rise less, thus generating such countercyclicality.

The first version of the New Keynesian model, as we describe it here, was presented by Woodford (1996); in Woodford (1999b), he introduced a definition of natural output also used here. A close precedent is Yun (1996), who looks at a similar setting but with monetary aggregates and without deriving an output gap-based New Keynesian Phillips curve or using an interest-rate rule. A paper that was also very influential in displaying the power of the New Keynesian setting is the review piece by Clarida et al. (1999) in the *Journal of Economic Literature*.

[27] Kiyotaki (1985) developed a model of monopolistic competition and identified an aggregate demand externality.

4.1 The New Keynesian model, brief description

As explained above, the New Keynesian framework arose as a result of trying to resurrect a framework describing output as demand-determined, where there is free capacity to produce more, so that higher demand for goods and services will be accommodated by supply, and where monetary policy can have an important role to play. The RBC model, in contrast, could perhaps be described as supply-driven, though in its construction it has both demand and supply, with perfect competition and flexible prices ensuring that supply equal demand at all times. Woodford's approach was not to construct an entirely different setting but instead build straight onto the RBC model, with the key change being an assumption that prices are "sticky", i.e., they do not adjust to clear markets. This also necessitated a framework with some degree of monopoly power so that the notion of "setting" a price could be given meaning.

The three-equation model describes a linear, stochastic, and dynamic system in three variables: the *output gap*, inflation, and the nominal interest rate. The output gap, a notion Woodford found to be the key variable for describing the evolution of output, is actual output minus the so-called *natural output* level. In his 2003 book, Woodford defined natural output to be that which would have arisen in the absence of price stickiness. Given that the framework is built on the RBC model, natural output could equivalently be described as the output that would result from the underlying RBC model.[28] This natural output definition is operationalized in a new way relative to the literature, and it is also not the same as another, common concept: full-capacity output. Here, as in the RBC model, output is determined by the work hours of households, which are variable, and a choice, on the individual level, implying that there is no useful notion of full capacity. The simplest model has no capital stock or investment goods, government expenditures are set at zero, and the economy is closed: net exports are zero.

The first model equation is the New Keynesian IS curve:

$$\tilde{y}_t = E_t\tilde{y}_{t+1} - \frac{1}{\sigma}\left(i_t - E_t\pi_{t+1} - r_t^n\right). \quad \text{(New Keynesian IS curve)} \tag{14}$$

Here, the variable $\tilde{y}$ denotes $y - y^n$, where y is (the log of) actual output and y^n is (the log of) natural output: $\tilde{y}_t$ is thus our output gap at time t (e.g., year or month). On the right-hand side, we have the output gap expected at t ("E_t") to prevail next period, at $t+1$ minus a constant multiplied by the expected real interest rate gap: the nominal interest rate i_t (between t and $t+1$) minus expected inflation minus the expected natural real rate, r_t^n, between t and $t+1$. The natural rate is the one that, again, would prevail absent price stickiness; Knut Wicksell, an inspiration behind Woodford's book, defined a similar natural interest rate in 1898.[29]

The New Keynesian IS curve, as we shall see below, is derived from the optimal consumption

[28] Here, "the RBC model" refers to an RBC model with monopolistic, not perfect, competition.
[29] Wicksell's book appeared in German: Wicksell (1898). Its translation to English appeared in 1936: Wicksell (1936).

behavior of households, taking into account that consumption equals output when supply adjusts to meet demand. It builds directly on Hall's consumption theory of consumption smoothing. Hence, absent interest-rate changes, changes in the marginal utility of consumption are unpredictable. A higher expected real interest rate will lower consumption demand today, an effect that is regulated by the curvature of utility with respect to consumption, σ: the higher the curvature, the less consumption will change.

In sum, whereas the traditional IS curve simply has real output as a negative function of the nominal interest rate, the New Keynesian IS curve has a similar dependence of current output on the nominal interest rate, but it is a "dynamic gap relation," relating how the output gap evolves over time to an interest rate gap. Moreover, the traditional IS curve only has the nominal rate, whereas expected inflation, an endogenous variable, enters here.

The second equation in the three-equation system is a forward-looking Phillips-curve:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa \tilde{y}_t \quad \text{(New Keynesian Phillips curve)} \tag{15}$$

This equation takes the place of the traditional Phillips curve. The traditional curve relates inflation negatively to the unemployment rate. Here, unemployment is replaced by the output gap: a higher output gap means higher hours worked relative to their "natural" level. Inflation between $t-1$ and t, π_t, depends positively on the output gap at t, and thus if the gap becomes negative (which can be seen as a parallel to having higher unemployment), inflation falls.

The New Keynesian Phillips curve is derived from the firms' price setting behavior and is fundamentally forward-looking. If other firms are expected to raise their prices in the future, a given firm will change its price now if it is able to, but not one for one: $\beta < 1$ in front of $E_t \pi_{t+1}$ is a discount factor. The current output gap also affects inflation because a higher output implies a higher marginal cost of production—the more actual output is raised above the natural level, the higher are wages, since workers need compensation for working more.[30]

The New Keynesian IS and Phillips curves are the key model equations: they show how the nominal interest rate affects output and inflation over time. What is missing is how this interest rate is set. Woodford took the view that modern central banks had moved away from primarily controlling the growth of the money stock. He also realized that it is possible to do away with money entirely in the model, without losing its usefulness for monetary stabilization. Hence, he added a rule whereby the monetary authority sets the interest rate to target an inflation rate by fighting departures from a target for inflation (which is zero in this, very simplest case), both by raising the rate when inflation rises and when demand pressure is high (the output gap is

[30]The "expectations-augmented" Phillips curve due to the 2006 Laureate in Economic Sciences, Edmund S. Phelps, allows the curve to shift; see Phelps (1967). However, it does not have the dynamic component that emerges naturally, as we shall see below, if a relationship has microeconomic underpinnings. The dynamic component, moreover, allows a more accurate description of how inflation behaves over time in response to shocks.

high):

$$i_t = \rho + \phi_\pi \pi_t + \phi_y \tilde{y}_t + v_t. \quad \text{(policy rule)} \tag{16}$$

Here, ϕ_π is greater than one: when inflation goes up by one unit, the nominal rate goes up by more than one unit, ensuring a rise in the real rate. This interest rate rule is often called the Taylor rule after his paper Taylor (1993). There, Taylor proposed this kind of rule as an appropriate one for the central bank to follow and showed that such a rule had matched the data quite well.

Finally, the policy rule contains a v_t: an unexpected, and exogenous, shock to the interest rate. This shock plays the role of shifting "demand" in the sense that a higher interest rate lowers consumer current demand (from the New Keynesian IS curve). The shock can be viewed as actions taken by the central bank that are hard to predict in advance by the market and, hence, surprises the economy. It is zero on average (in expectation), and when it is positive, it is contractionary.

Together, the three equations describe a complete dynamic (stochastic) system that can be solved and simulated for different parameter values and shock realizations. The system has two shocks: in addition to the demand (interest-rate) shock, the natural interest rate r_t^n is a shock variable too, representing supply shocks, such as those to productivity. In particular, a positive productivity shock today (relative to tomorrow) would amount to a lower natural real interest rate, since it would imply higher output, and hence make consumption today less scarce, relative to tomorrow: the real rate, conceptually, is the market value of goods today relative to tomorrow. But we see in the New Keynesian IS equation that $\tilde{y}$ would actually fall, unlike y, if the natural real rate falls. That is because although actual output rises, it does so less than it would absent price rigidities.

Hence, from the policy rule, the central bank's reaction is to lower its rate, thus encouraging the economy to raise its output faster, toward where it would be absent price rigidities. In contrast, a positive demand shock (a lower interest rate than expected), by stimulating output, hence generating upward pressure on prices, would cause the central bank to immediately "lean against the wind" and raise the interest rate: a higher inflation on the right-hand side of the policy rule makes the central bank want to raise rates to slow the economy down. Thus, positive demand and supply shocks cause the central bank to move its instruments in different directions (up and down, respectively).

Based on the New Keynesian model, in his 2003 book Woodford also provides support for Taylor's assertions: the model implies that optimal policy—as defined by the maximization of consumer welfare—can be pursued with this kind of rule. That is, the responses to the shocks described above are also correct, qualitatively, from a normative perspective. This stands in contrast to the traditional Keynesian framework, which does not come with an explicit norma-

tive criterion, as that model's equations are not derived from, and thus not tied explicitly to, consumer welfare.

Another key feature of the three-equation system is that it can be used to evaluate counterfactuals: how output and inflation would behave for different alternative interest rate choices by the central bank, or across different economies that differ in their New Keynesian IS- or Phillips-curve coefficients.[31]

The three-equation system does, of course, also leave out many important variables and mechanisms (such as investment or wage stickiness), which are important in practice. However, the New Keynesian model has shown to be easy to extend into larger systems with many more variables. In particular, additional demand and supply shocks can be included. These extensions all build directly onto the simple three-equation model. Many of them were presented in Woodford's 2003 book; we discuss some of them briefly in Section 4.5 below.

4.2 The underlying microeconomics

In this section, we describe the microeconomic foundations behind the three-equation New Keynesian model. The real economy underlying the model is actually static, as the model does not have physical capital: all the internal dynamics are due to prices being sticky, i.e., hard to adjust over time.[32] We describe the different model components in turn, beginning with the consumer.

4.2.1 Consumers: enter Hall

We assume that there is a representative household that chooses stochastic sequences for consumption, hours worked, and saving in bonds to maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u\left(C_t, H_t\right),$$

subject to a budget in nominal terms, reading

$$P_t C_t + Q_t B_{t+1} = B_t + W_t H_t + \Pi_t.$$

This is a version of Hall's model, extended to include a choice for hours worked, H_t, in each time period, and with a more complicated-looking budget constraint, since it is written in nominal

[31] A qualification is necessary here: not all policy counterfactuals generate a sharp prediction, in the sense of a unique solution to the dynamic three-equation system. For example, as will be briefly touched on in Section 4.4 below, a value of $\phi_\pi < 1$ can be consistent with self-fulfilling inflationary paths as well as stable inflation; hence, the model equations do not deliver a sharp answer.

[32] Movements over time, of course, also occur due to movements in exogenous variables, such as productivity.

terms. P_t is the nominal price level in period t, and we will define inflation between periods t and $t+1$ as $\pi_{t+1} \equiv (P_{t+1} - P_t)/P_t$. C_t is a real variable, as we will explain below, but unlike in the description of Hall's work above, it is actually a bundle of different underlying goods that the consumer buys; hence, P_t is a (nominal) price index for this goods bundle. Moreover, saving occurs through nominal bonds: Q_t is the price of nominally riskless one-period bonds B_{t+1} paying out 1 currency unit at $t+1$.[33] Thus, $1/Q_t - 1 \equiv i_t$ is the nominal interest rate between t and $t+1$. W_t is the nominal wage and $W_t H_t$ nominal labor earnings; Π_t is additional, exogenous income, the source of which we will describe below. Thus, comparing to (3), we have in essence the same dynamic problem, except here there is an additional choice (of hours worked).

In practical applications, the utility function is specified as $u(C,H) = \frac{C^{1-\sigma}-1}{1-\sigma} - \chi \frac{H^{1+\frac{1}{\psi}}}{1+\frac{1}{\psi}}$. Taking first-order conditions for C_t and H_t, we have, for all t,

$$\begin{aligned} C_t^{-\sigma} &= \beta \, \mathbb{E}_t \left[\frac{1+i_t}{1+\pi_t} C_{t+1}^{-\sigma} \right] \\ C_t^{-\sigma} \frac{W_t}{P_t} &= \chi H_t^{\frac{1}{\psi}}. \end{aligned}$$

The first of these equations is Hall's equation (8) with some further manipulation, to be developed into the New Keynesian IS curve.[34]

4.2.2 Firms and sticky prices

Firms sell differentiated products, as modeled by Dixit and Stiglitz (1977), and the demand for their products follows from the "love of variety" of consumers: the object delivering utility from consumption at t, C_t, is given by

$$C_t \equiv \left(\int_0^1 C_t(i)^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}},$$

where $C_t(i)$ is consumption of variety i. Total consumer spending is $\int_0^1 P_t(i) C_t(i) di$, where $P_t(i)$ is the price of variety i at t.

We can now derive a so-called exact price index. This is the cost per unit of C when the consumption basket is composed so that it minimizes the total cost for the consumer given all the underlying $P_t(i)$s. It is straightforward to show that this exact price index is given by

$$P_t = \left(\int_0^1 P_t(i)^{1-\varepsilon} di \right)^{\frac{1}{1-\varepsilon}}. \tag{17}$$

[33] Bonds are zero in equilibrium in the simplest version of the model. Although we omit them here for readability, there are also state-contingent bonds in the budget; their prices, which are set so that each state-contingent bond is exactly in zero demand in equilibrium, are used later.

[34] Set $\tau = 1$ in equation (8) and note that the real interest rate satisfies $1 + r = (1+i)/(1+\pi)$.

The price index depends on all the prices, but also on the elasticity of substitution ε between the different consumption goods. If this is high, a high price leads to the weight of that good in the consumption basket being low and thus less important for the price index than if ε is low. When ε approaches zero, all goods need to be consumed in the same proportion and the price index is simply the unweighted arithmetic average of all prices.

Based on consumer optimization, we can also derive the demand for each good i as

$$C_t(i) = C_t \cdot \left(\frac{P_t(i)}{P_t} \right)^{-\varepsilon}.$$

This constant-elasticity demand function is key for understanding the firm's problem.

The monopolistic competitors Each good i is produced by a monopolist with a production function that is, in its very simplest form, linear in labor: production of good i obeys $Y(i) = AH(i)$, where $H(i)$ is the amount of labor used to produce good i and A is labor productivity. The marginal cost of a unit of output is then W/A, where W is the wage. Suppose, first, that the firm is free to choose its price every time period. Then the firm would not face a dynamic choice problem—it faces a demand curve at each point in time so the overall profit maximization problem can be split up to profit maximization problems for each t. Thus, the firm that produces good i in period t solves:

$$\begin{aligned} & \max_{P_t(i), Y_t(i)} P_t(i) Y_t(i) - \frac{W_t}{A_t} Y_t(i) \\ \text{s.t. } Y_t(i) &= C_t \cdot \left(\frac{P_t(i)}{P_t} \right)^{-\varepsilon}. \end{aligned}$$

This problem, which is only well defined for $\varepsilon > 1$ (i.e., when there is sufficient substitutability across goods), is solved by

$$P_t(i) = \frac{\varepsilon}{\varepsilon - 1} \frac{W_t}{A_t}.$$

Firms thus set a constant markup $\mu \equiv \frac{\varepsilon}{\varepsilon - 1}$ over their marginal costs $\frac{W_t}{A_t}$ (which vary over time).

Natural output and the natural rate Recall that Woodford defined natural output as the output that would materialize in a market equilibrium in the absence of sticky prices. Such a market equilibrium is one with monopolistic competition where, due to the symmetry assumed here across firms, all firms make the same choices. Hence $P_t(i) = P_t(j) = P_t \; \forall i, j \in [0, 1]$. In the dynamic model with price stickiness, however, different goods will have different prices and hence be consumed in different quantities. This will mean a welfare loss for the consumer, as the assumptions on the utility function—strict concavity in the consumption of each type of good and symmetry—call for equal consumption.

When all prices are identical, $P_t = \mu \frac{W_t}{A_t}$, or $W_t/P_t = A_t/\mu$. Since $\mu > 1$ due to the pricing power, the (real) wage is below the the marginal product of labor A. This allows us to fully solve for a market equilibrium for each period t separately. We can use the household's first-order condition between consumption and leisure, $u_1(C_t, H_t)(W_t/P_t) = -u_2(C_t, H_t)$, the market-clearing condition $C_t = Y_t$, and the production function $Y_t = A_t H_t$ to solve for C_t, Y_t, and H_t. The obtained level of output is denoted by the natural level of output, Y_t^n (or y_t^n in logarithms). Its level will depend on A_t, labor productivity. A change in A_t can then be seen as a supply shock: a movement in productivity. This feature makes the model connect very closely with the RBC model developed by Kydland and Prescott (1982), for which the main source of fluctuations is the A_t.

We can also compute the expected real natural rate of interest, r_t^n, using Hall's Euler equation for consumption: we just have to insert the obtained C_t and C_{t+1}.[35] Note that Woodford's concept of a natural rate is determined without reference to nominal variables. The Wicksellian concept can thus be derived in an entirely real economy; departures from the "natural real allocation" are key in driving inflation off its target.

Sticky prices We now describe the key friction in the model: costs of price setting. With this friction, the firm's problem becomes dynamic and there is a role for the central bank. The price level, denoted in monetary units, will in general vary over time, as the economy is subjected to shocks.

Two common settings for describing price-setting frictions are used in applied work:

1. **"Rotemberg pricing"**: Rotemberg (1982b) assumed that there is a resource cost associated with changing from price $P_{t-1}(i)$ to $P_t(i)$ that is equal to $g(P_t(i)/P_{t-1}(i))$. The function g is assumed to be quadratic with a minimum at $g(1) = 0$, i.e., it is costless to maintain constant prices.

2. **"Calvo pricing"**: Calvo (1983) assumed an extreme form of menu costs where, each period, a firm can adjust its price with probability $1 - \theta$ ("the firm is visited by the Calvo fairy"). In other words, there is a cost of changing prices that is either 0 (when the fairy visits), or else it is ∞. The opportunity to reset prices is thus a stochastic variable and it is assumed to be independent across firms and time.

In the literature, both formulations are common, and a range of other price-setting assumptions have been used as well. Woodford's most commonly used main three-equation model builds on Calvo pricing, so we will adopt it in our description here.

[35] With some algebra, we obtain, for output, $Y_t^n = (\chi\mu)^{-\frac{\psi}{\psi\sigma+1}} A_t^{\frac{\psi+1}{\psi\sigma+1}}$. The natural real rate of interest between t and $t+1$ is random and denoted $r_t^n(A_{t+1})$. It can be computed by applying the Euler equation state by state: $1 = \beta(1 + r_t^n(A_{t+1}))(A_{t+1}/A_t)^{-\sigma\frac{\psi+1}{\psi\sigma+1}}$. The expected natural rate, r_t^n, then equals $E_t(r_t^n(A_{t+1}))$.

Under Calvo pricing, the probability that a given firm resets its price today is $1-\theta$, that it does not reset until tomorrow is $\theta(1-\theta)$, the day after $\theta^2(1-\theta)$, and so on. Given this structure, it is straightforward to derive a recursive formula for the dynamic evolution of the price index:

$$P_t^{1-\varepsilon} = (1-\theta)(P_t^*)^{1-\varepsilon} + \theta P_{t-1}^{1-\varepsilon}. \tag{18}$$

The equation tells us that the price index today (raised to a power) is a weighted average of the price set by all the firms that can reset today, P_t^*, and the previous period's price index (each raised to the same power).[36] We now turn to the key firm decision, which is a dynamic one: how P_t^* is determined.

The forward-looking firm's problem Under Calvo pricing, a firm i that can choose its nominal price at time t, $P_t(i)$, realizes that the output it produces at all times after that in case the price cannot be changed until then—$Y_{t+s|t}(i)$ for all $s \geq 0$—will be determined by demand. Thus, the firm maximizes

$$E_t \sum_{s=0}^{\infty} \theta^s Q_{t,t+s} \left(P_t(i) Y_{t+s|t}(i) - W_{t+s} \frac{Y_{t+s|t}(i)}{A_{t+s}} \right) \tag{19}$$
$$\text{subject to} \quad Y_{t+s|t}(i) = C_{t+s} \left(\frac{P_t(i)}{P_{t+s}} \right)^{-\varepsilon}.$$

Here, the choice variables are $P_t(i)$ and $Y_{t+s|t}(i)$ for all $s \geq 0$. The objective is the discounted sum of profits at all future times in the event the firm is bound by the price it chooses today.[37] For a time s periods ahead, this occurs with probability θ^s. The constraint is the demand faced by the firm at $t+s$ in that event. In addition to its own price, it depends on aggregate demand C_{t+s} and the aggregate price level P_{t+s}.

4.2.3 Closing and solving the model

The specifications above for consumers and firms, along with market clearing—that is, $Y_t(i) = C_t(i)\ \forall(i,t)$ and $H_t = \int_0^1 H_t(i)\, di$—complete the model description, except for the addition of a specification of how interest rates are set. Here, the Taylor rule is a standard choice and we discuss it briefly below.

How can the model be solved? As usual, when it comes to dynamic models with rational expectations, the solution is a fixed point in stochastic sequence space. The reader will also note that most of the equations are nonlinear; hence, the solution of the model will need to rely

[36] This formula is a straightforward application of the Calvo structure and the functional form of the index.

[37] The (time- and state-dependent) nominal discounting from t to $t+s$ comes from the Euler equation applied state by state and reads $Q_{t,t+s} = \beta^s \left(\frac{C_{t+s}}{C_t}\right)^{-\sigma} \frac{P_t}{P_{t+s}}$.

on numerical methods. A key ingredient to Woodford's approach is to approximate the dynamic system with linearization methods around a steady state, where A_t is constant and $v_t = 0$. In the simplest version of the setting as shown in the specification above, the inflation target is zero, and hence, the steady state has constant and equal prices for all firms. If a shock then hits, it will lead to dynamics that converge back to steady state, and Woodford shows that these dynamics are precisely described by the three-equation system; if recurrent shocks materialize, the system will fluctuate around this steady state.

We will not detail the full linearization, but we will illustrate with some key steps, focusing on how to derive the New Keynesian IS equation. The Euler equation, in the absence of uncertainty, reads as follows.

$$\begin{aligned} Y_t^{-\sigma} &= \beta \frac{1+i_t}{1+\pi_{t+1}} Y_{t+1}^{-\sigma} \\ \implies & \left(\frac{Y_t}{Y_t^n}\right)^{-\sigma} = \beta \frac{1+i_t}{1+\pi_{t+1}} \left(\frac{Y_{t+1}}{Y_{t+1}^n}\right)^{-\sigma} \left(\frac{Y_{t+1}^n}{Y_t^n}\right)^{-\sigma}. \end{aligned}$$

Using the definition $\tilde{y} \equiv \frac{Y_t}{Y_t^n} - 1$ and recalling that the natural real interest rate is the equilibrium interest rate when output is at its natural level—here under perfect foresight, in the full model in its expected value—implying that $\frac{1}{1+r_t^n} = \beta \left(\frac{Y_{t+1}^n}{Y_t^n}\right)^{-\sigma}$, we have

$$(1+\tilde{y}_t)^{-\sigma} = (1+\tilde{y}_{t+1})^{-\sigma} \frac{1+i_t}{1+\pi_{t+1}} \frac{1}{1+r_t^n}.$$

We note the power of Woodford's definition of natural output and the natural real rate. By taking logarithms and first-order approximating around the steady state, where $\tilde{y}_t = 0$, we obtain

$$\tilde{y}_t = \tilde{y}_{t+1} - \frac{1}{\sigma}\left(i_t - \pi_{t+1} - r_t^n\right).$$

Apart from lacking an expectations operator, we now have an equation that reproduces the New Keynesian IS curve, (14).[38] Notice that the supply shock (changes in the A parameter) does not directly affect our output variable here, since output is expressed as a gap: if actual output follows natural output, the productivity variable will not have consequences. However, it enters—as we noted above—through the appearance of the natural real interest rate, r_t^n.

The derivation of the New Keynesian Phillips curve, (15), follows a conceptually identical, but significantly longer, derivation using the firm's first-order condition for price-setting in its problem (19). This first-order condition involves an infinite number of terms, but these terms

[38] To obtain the version with an expectations operator, just apply the Euler equation state by state, perform the same linearization, and then add across states weighted by the respective probabilities. This procedure illustrates "certainty equivalence": studying the deterministic dynamics of a linear system suffices for characterizing the system with random shocks.

have a recursive structure. The recursive structure, which determines P_t^*, can in turn be combined with the dynamic formula for the price index, (18), and can be log-linearized to yield the New Keynesian Phillips curve. Inflation today becomes an increasing (linear, here) function of inflation in the future, representing how firms are forward-looking: a firm that can reset the price today raises the price more, and proportionately so, if other firms raise their prices in the future, to keep up with competition. Future inflation is downweighted, however, since the effect of future price competition on firms' present-value profits is discounted, using the market discount factor, $1/(1+r^n)$, which equals β. Inflation is also increasing in production: when output rises, marginal costs rise because wages rise: workers have to work more hours and require compensation for that.[39]

The summary effect of output on inflation in the New Keynesian Phillips curve is given by the parameter κ, often referred to as the slope of the Phillips curve. A big κ means a high slope: inflation reacts strongly to output; conversely, a higher inflation rate barely moves output. Of course, κ is a function of the deep parameters of the model: it can be shown to equal $(1-\theta)(1-\beta\theta)(1/\psi+\sigma)/\theta$. The expression reveals that the slope is larger if prices are more flexible (a lower θ) or if marginal cost is more sensitive to the level of production ($1/\psi + \sigma$ is high, reflecting sharply increasing marginal costs of working relative to consuming).

We now have the two main model equations; the last equation, the policy rule, completes the dynamic system and, as we shall briefly discuss below, ensures that the solution to the system is unique. The addition of uncertainty is straightforward, and hence we have the three-equation model as described above: equations (14)–(16), where it should be noted that ρ in the interest rate rule equals $1/\beta - 1$: the steady-state real rate r^n.

4.3 Illustration of the workings of the three-equation model

The simple model allows us to interpret business-cycle movements as a result of recurring disturbances, in the form of supply and demand shocks, and the economy's response to them: the propagation mechanism. We will now consider a few simple examples to illustrate these elements. The discussion will begin by describing how to formally analyze two simple examples and then show a richer example; finally, we will discuss how appropriate stabilization policy may in some case involve reinforcing an output contraction.

Two fully solved simple examples Observe that in the three equations, no variable—neither i_t, nor $\tilde{y}_t$, nor π_t—is predetermined. In the absence of shocks, the equation system can then

[39] The three-equation model is often used with a production function that is not linear in labor, but strictly concave. This is then a further source of rising marginal costs as output expands.

be solved directly: $\tilde{y}_t = \pi_t = 0$ and $i_t = \rho$ at all times.[40] Suppose, however, that at t we obtain an unexpected $v_t > 0$: what are the dynamics that follow? For simplicity, suppose v is zero in the future. Then at $t+1$, since no further shock appears, $\tilde{y}_{t+1} = \pi_{t+1} = 0$ and $i_{t+1} = \rho$, and the same goes for all future periods. What remains is to solve for $\tilde{y}_t$, π_t, and i_t; it is easy to see, by eliminating π_t and i_t from their respective equations, that $\tilde{y}_t = -(1/\sigma)(\phi_\pi \kappa \tilde{y}_t + \phi_y \tilde{y}_t + v_t)$. Clearly, an upward shock to the nominal rate contracts output, and as a result, inflation falls.[41] Thus, we have real effects of monetary policy.

A second experiment comes from a sudden rise, away from a steady state, in productivity, A_t. Again, assume this is a one-time shock and that all future values of A are at their steady-state value. This shock will, as mentioned above, affect the economy by changing the real rate, r^n. Recall from above that the natural (gross) real rate equals $\frac{1}{\beta}\left(\frac{Y^n_{t+1}}{Y^n_t}\right)^\sigma$, where Y^n_{t+1} is back at steady state and where Y^n_t can be readily solved for from the static model. The static model, quite naturally, implies an increase in Y^n_t as A_t rises, in part through an effect on the endogenously determined hours worked, H_t. Thus, the impact effect from (14) is that the natural real rate will fall, implying that the output gap thus falls, perhaps surprisingly.[42] Intuitively, natural output rises, but recall that the model with sticky prices implies that many firms cannot lower their prices to generate higher demand; they would if they could, since marginal costs have fallen, and thus actual output rises less than does natural output. Hence, the gap decreases, and the supply shock is de facto weakened by the stickiness of prices.

Persistent shocks How the economy responds to each kind of shock—the initial impact and how it then propagates through the economy over time—in part depends on its own inherent dynamics. For example, suppose the supply shock is an increase in energy prices, which we can think of as a lower productivity, A.[43] Moreover, say the energy price hike is expected to be persistent but slowly disappear over time. The firms that do adjust prices immediately will raise their prices, and more so to the extent the higher energy price will be expected to last. The average behavior of the entire economy will then also depend on how strongly consumers respond to the implied increased real rate ($1/\sigma$ in the IS curve), which in turn is also a function of how strongly the central bank reacts to the developments (the ϕs). Unlike in the two examples above, the expected future behaviors of energy prices and the economy's reaction to them will affect outcomes today. This forward-looking feature was entirely absent in the traditional Keynesian

[40] In the background, the price level of those firms not able to adjust their price is predetermined, but in steady state, they are all the same and the level of prices can be regarded as a numéraire.

[41] In applications, v_t is often viewed to be positively serially correlated, and a one-time rise is followed by a contraction that is dampened over time. Then the model is less easy to solve, but standard methods for linear difference equations can readily be applied.

[42] Of course, there is a reaction of inflation as well as the nominal rate, but these will not overcome the impact effect.

[43] Energy is not an explicit production input in our equations above, but it is easy to incorporate it, in which case A will be a negative function of the energy price.

framework.

Stabilization policy The forward-looking nature of the New Keynesian model also suggested new avenues to pursue for policymakers. If, in particular, a central bank announces a path for future nominal interest rates, it would affect the economy today. Such *forward guidance* was proposed by Woodford and others as an additional means to stimulate the economy when current nominal interest rates cannot be lowered further. This method was adopted during the zero lower bound (ZLB) period when rates had reached zero (and in some cases had even become negative); in the US, the ZLB period is usually dated 2008–2015.[44] At such a time, signaling that future nominal rates will remain very low would hold inflation back and hence lower the real rates, thus also stimulating current demand. Of course, forward guidance requires credibility of the central bank's announced path, which is not to be taken for granted.

A second insight the framework offers is that the output gap itself directly reveals the market failure—the sluggish reactions of firms to price changes—appearing in response to shocks. The energy price hike discussed above makes actual output fall, but natural output falls even more: the gap becomes positive. The response of the central bank, since $\phi_y > 0$, is to raise rates to slow the economy down further. This response makes sense: even though output falls, it is appropriate to push it down even more, given that firms on average do not fully price in the increased marginal costs. Thus, central bank actions that may appear destabilizing can be appropriate and reflect sensible stabilization policy. Of course, the nature of the shock is key to whether a fall in output should be counteracted or reinforced.

4.4 Discussion of assumptions

We now discuss some important assumptions underlying the analysis, along with a number of extensions of it.

Cashlessness The three-equation model is very simple and easy to use in illustrations. The simplicity, however, can also be deceiving; the reader may for example wonder how there can be a concept of "nominal" without money appearing in any of the equations. On the one hand, this means that a nominal money stock need not be essential to the key workings of monetary policy. In fact, the role of money in the model is simply as the unit of account in which prices are assumed to be sticky. On the other hand, Woodford (1998) also showed that, under some conditions, the three-equation system can be derived as the limit of an economy where a money stock exists and is explicitly valued, but where the extent to which it is valued goes to zero. This can be

[44] Different central banks appeared to consider different value as their lower bound, i.e., a rate they were uncomfortable going below; there were fears that sufficiently low deposit rates would cause mass withdrawals of funds.

viewed as a description of current-day money use in some economies, like Sweden's, where the demand for currency by consumers is very close to zero. Woodford's demonstration is rather subtle and is intimately connected with equilibrium indeterminacy and the use of interest-rate rules, so let us briefly also comment on these issues.

Potential indeterminacy and the interest-rate rule In the Keynesian model as described in traditional undergraduate textbooks, monetary policy is not determined by an interest rate rule but as conducted by open-market operations, with the result that we obtain an LM curve linking the nominal interest rate to the stock of money, controlled by the central bank, and to real output.[45] As explained above, this derivation hinges on a relationship describing the real demand for money, M/P, as a function $L(Y, i)$, where Y is real output and raises the need for real balances and i is the nominal interest rate, which lowers money demand when it rises, as money bears no interest but other assets do. If the central bank were to simply set i, then an indeterminacy appears: M/P is determined but M and P are not pinned down separately. In other words, one can imagine an unstable price level arising from a central bank simply choosing an interest rate.

This problem was discussed in the literature prior to Woodford's work and regarded as an advantage for rules involving monetary targets, where this kind of indeterminacy is not present; see Sargent and Wallace (1975). Woodford's analysis, however, amounted to a defense of the use of interest-rate rules: when the rate is set in such a way that it steers sufficiently strongly toward a target, such as the price level or the inflation rate, the indeterminacy disappears. Recall that ϕ_π is the response in the Taylor rule for the central bank to changes in the inflation. $\phi_\pi > 1$ implies that increased inflation leads to increases in the nominal interest rate larger than one-for-one. This means that higher inflation tends to lead to higher real interest rates.[46] This was very useful also because central banks perceive themselves as precisely setting the rate, but also setting it purposefully to prevent inflation hikes. Thus, the result both provided an accurate description of how monetary policy is conducted and how it should be conducted to provide stability of expectations. Woodford (1998) precisely relied on an interest-rate rule in deriving the cashless limit: M/P goes to zero by M approaching zero, while maintaining a finite P.

Woodford's main way of resolving indeterminacy—the use of an interest-rate rule with high enough values for ϕ_π and ϕ_y—is not necessarily the only available possibility. Some authors,

[45] Since Woodford's 2003 book and its impact on academia and on policy, the interest-rate rule has started to replace the money stock as the main instrument in textbooks.

[46] This condition ensures that, locally around the steady state, there is a unique path leading back to the steady state. It does not rule out locally exploding paths, i.e., paths that lead elsewhere in the long run. Benhabib et al. (2001a) and Benhabib et al. (2001b) show that such paths can be equilibria under some conditions and propose ways to ensure global uniqueness. In Chapter 2 of his 2003 book, Woodford discusses these issues in detail. For a related discussion of indeterminacy, its causes, and other refinement conditions in models where output is demand-determined, see Angeletos et al. (2026).

including Woodford himself, as in Woodford (1995), have discussed the so-called fiscal theory of the price level as one complement, or alternative. Such an approach does hinge on having nominal assets and what is referred to as "fiscal dominance": the price level, while being targeted by the central bank, is also determined by market expectations that the intertemporal budget constraint for the government must hold exactly. In concrete terms, suppose the fiscal authority raises spending without any plans to finance it with taxes. One possibility is then that the price level rises so as to devalue, in real terms, the value of nominal government debt, which then makes the higher spending budget feasible. This alternative theory, which relies on market participants to coordinate on the price-level rise, is preferred by some academics, but to many others does not reflect how fiscal and monetary policies interact.

4.5 Extensions, applications, and challenges

As the discussion above reveals, the analysis behind each of the assumptions contains important insights, and these may be overlooked when studying the highly stylized three-equation model. These assumptions are also important, as their validity remains relevant when looking at less stylized versions of the New Keynesian model. We now briefly discuss some important extensions, an example of how the New Keynesian analysis is used in practice, and some challenges.

Extensions The literature applying and extending the New Keynesian framework is very large, both in terms of theoretical developments and applications. One line of work addresses aggregate data and introduces more macroeconomic variables and more drivers of shocks; we very briefly outline one such model below. Conceptual, theoretical extensions here—which are also important for understanding the data—involve the introduction of sticky wages as well as unemployment into the framework.

When it comes to assessing how well New Keynesian models fit the data, the most well-known early papers include Smets and Wouters (2003, 2007), looking at euro area and US data, respectively, and Christiano et al. (2005), again relative to US data. These papers structurally estimate "medium-scale" New Keynesian models, i.e., models with more shocks and variables.[47]

Earlier, Rotemberg and Woodford (1997) estimated a New Keynesian model, but a simpler version than those that later came to dominate the applied literature. Other theoretical extensions include a more flexible preference structure, whereby the real markups are allowed to fluctuate cyclically, thus adding a second mechanism to generate countercyclicality in mea-

[47] The Smets and Wouters frameworks use the following additional features: (i) sticky wages; (ii) wage and price indexation (motivated by practice); (iii) capacity utilization costs; (iv) adjustment costs to investment; (v) consumption habits; and (vi) more shocks—such as investment-specific shocks, government expenditure shocks, risk-premium shocks, and price and wage markup shocks—amounting to a total of 7 shocks, including productivity and monetary policy shocks.

sured markups.[48] The basic model has also been extended to include rich heterogeneity, e.g., sectoral linkages, intermediate goods and input-output linkages, as the degree of price stickiness varies across goods; for example, see Rubbo (2023).

Another important line of work, an example of which is discussed below, is the extension to an open economy and the analysis of the exchange rate volatility and the transmission of monetary policy across economies. This literature includes analysis of "small" open economies, as in Clarida et al. (2001) and Galí and Monacelli (2005), thus creating New Keynesian versions of the Mundell-Fleming model; and frameworks with interactions across larger regions, such as in Corsetti and Pesenti (2001) and Corsetti et al. (2008).[49]

A more recent, and very vibrant, literature builds in consumer heterogeneity by incorporating uninsurability of idiosyncratic shocks to the incomes of different consumers. This so-called HANK literature (HA for Heterogeneous Agents) connects more directly to the theoretical and empirical literature generated by Hall and Mishkin (1982), and it delivers a framework where the marginal propensities to consume differ widely between consumers.[50] An even more recent literature studies firm heterogeneity in the context of the New Keynesian model; here, the marginal propensity to invest differs across firms.

The analysis above takes as given an interest-rate rule of a certain kind and is chiefly descriptive. Another conceptual issue, then, is that of optimal monetary policy: given that the world can be described by a model of the New Keynesian variety, what would the appropriate way of setting interest rates (or a monetary quantity) be? The literature on the best conduct of monetary policy is vast, both prior to and after the New Keynesian model emerged, and it is impossible to do it full justice here, but a few points are worth noting.

In their 1997 paper, Rotemberg and Woodford (1997) looked at the normative performance of different monetary rules given their estimated New Keynesian model—since such a model has a well-defined notion of economic welfare (given by the present-value utility of the representative agent). The same year, Goodfriend and King (1997) made the conceptual point that the goal of monetary policy in the New Keynesian model can be viewed as trying to move the economy back toward the RBC allocation that would result in the absence of price stickiness: that is, try to "undo" the stickiness.[51]

For the simple three-equation model, a second point is that if we replace the intercept in the interest-rate rule, ρ, with r_t^n, the expected natural real rate, then it is possible to maintain price

[48] In the basic model, price over marginal cost moves countercyclically: as marginal costs rise when output expands, price sluggishness makes the ratio fall. In extensions based on Kimball (1995), consumers are assumed to exhibit lower sensitivity to price as their consumption of a good rises; this channel allows the countercyclicality to be high even for moderate levels of stickiness.

[49] Robert A. Mundell was awarded the Prize in Economic Sciences in 1999.

[50] In some of these models, central bank policy can influence the long-run real rate, a feature that is not present in the basic model but whose relevance is debated.

[51] The basic model also has a distortion due to monopoly power, which is arguably better dealt with using regulatory or other policy instruments.

stability (zero inflation) and close the output gap entirely, no matter how the technology shock moves.[52] This phenomenon was dubbed the "divine coincidence" in Blanchard and Galí (2007), who, however, argued that the New Keynesian model will lose this feature when real wages are also sticky. Then, monetary policy involves a "second-best" tradeoff between maintaining inflation close to target and making the output gap smaller.[53] Hence, there is extensive follow-up work, by Woodford and others, discussing the properties of different rules. Woodford (2003) showed, in particular, how a rule similar to that in (16) is optimal under some conditions.

A third point is the challenge of time inconsistency, the notion brought up in Kydland and Prescott (1977) and discussed in Barro and Gordon (1983): in general, the optimal conduct of monetary policy would require commitment to future policy choices, as there will be incentives to deviate from a policy plan that is optimal today. Woodford (1999a) suggested that a time-independent rule could nevertheless be appropriate from what he labeled a timeless perspective: the commitment plan can be characterized as converging over time to a time-independent rule.

In the above description of straightforward extensions of the model, a major extension is hidden: when government expenditures, taxes and transfers, along with government debt, are included, the model can be applied directly for fiscal policy evaluation. In fact, many government agencies as well as international organizations around the world do use such frameworks, in each case calibrated to their particular circumstances.[54] To what extent could, for example, a lower tax rate, or government transfers, help the economy out of a recession, through the boost on demand that it would generate? For such questions, a model where output is demand-determined is required, and the New Keynesian model is now *the* model that is used for answering such questions. Clearly, the subject of how fiscal policy can be used as a stabilization tool is a major one for all of economics. The need to address this topic is also an important reason behind the impact of New Keynesian models. In this document, we focus on monetary aspects, as they were the key focus for the researchers who developed the model.

It must also be mentioned that in parallel with the development of the New Keynesian model, there have been major advances in computational methods, allowing researchers to solve the more applied, larger-scale models with ease. One of the key developments was the construction of DYNARE, a convenient package that has become a core tool both for the nu-

[52] Let v_t equal zero and note that the three equations are then solved by $\pi_t = 0$ and $\tilde{y}_t = 0$ for all t. This feature holds for any values of ϕ_π and ϕ_y, and the assumption of high enough values for these parameters then ensures local determinacy.

[53] In fact, the baseline model Woodford (2003) used to study policy has three shocks, the additional shock being a disturbance in the price-setting equation that additionally moves the economy away from its natural level. The presence of this so-called cost-push shock, which is given microeconomic underpinnings in later research, also implies a non-trivial policy tradeoff.

[54] In Sweden, a fiscal-policy model—labeled SELMA—has been constructed at the National Institute of Economic Research. Some agencies, such as the International Monetary Fund, of course also use rich models to jointly analyze fiscal and monetary policy.

merical solution and econometric estimation of macroeconomic models, and one of its main early uses was precisely to extend and analyze the New Keynesian model.[55]

An application: MAJA—the New Keynesian model developed at Sveriges Riksbank We now give a description of a concrete extension of the basic New Keynesian model to a single economy—Sweden's. MAJA (*Modell för Allmän JämviktsAnalys*—Model for General Equilibrium Analysis) is the latest model used by Sveriges Riksbank for the interpretation of economic developments, forecasting, scenarios, and policy evaluation. It was developed by Corbo and Strid (2020) and builds on the work of Christiano et al. (2005) as well as on Smets and Wouters (2003). Like those settings, it extends the New Keynesian framework in various dimensions and is an example of its usefulness and adaptability. Similar extensions, designed to fit national or regional specificities, exist and are in active use at central banks and finance ministries all over the developed world.

A key extension of MAJA relative to the model described above is the introduction of foreign trade. Sweden is a small open economy where around half of demand comes from exports. Sweden is assumed to be small, in the sense that it cannot affect the world economy. The domestic Swedish economy is taken to consist of a number of different firms, which produce, import, or export different types of goods; households, which obtain utility from consumption and leisure and offer their labor services to firms; a government; and a central bank. Each firm produces a differentiated domestic intermediate good, using both labor and capital and produces to meet demand. Price stickiness is introduced by using the Calvo assumption. The result is a New Keynesian Phillips curve with similar properties to the one discussed above.

Wages are set by (small) unions representing different types of labor, rather than assuming a competitive labor market. These unions then choose to restrict employment in order to maintain higher wages, implying unemployment at the chosen wages.[56] As labor is monopolistically supplied, wages are set with a positive markup, which generally results in wages that are too high to ensure full employment. Just as the monopolistic firms setting prices, unions are assumed to set wages in a staggered fashion subject to Calvo wage setting frictions. Once the wage has been set, households meet the demand for workers at the going wage. In every period, each union faces a fixed probability that it can re-optimize its nominal wage. This leads to a separate New Keynesian Phillips curve for wage inflation.

In MAJA, the monetary policy set by the central bank is represented as a slightly modified interest-rate rule, which adjusts somewhat more sluggishly by the inclusion of lagged policy rates. A government that determines public spending is also included. MAJA is substantially larger than the Smets-Wouters models: the core model has 29 exogenous shocks, of which 16

[55] The original DYNARE code was developed by Michel Juillard in 1994.
[56] Some New Keynesian models instead use search matching theory to introduce unemployment into the model.

are domestic. These include shocks to productivity, to the price and wage setting powers, to preferences, to energy prices, to the exchange rate, and to government consumption.

The model can be used to account for the historical importance of the different shocks for inflation and other endogenous variables. An example is given in Figure 1 below, which decomposes Swedish inflation over the period 1996–2019 into the contribution from different types of shocks. Different shocks have been key during different periods, but over the entire period covered, we see a downward trend in inflation, as revealed by the blue area, even as the trend has also fluctuated somewhat. As for the factors behind short-run fluctuations, we note that energy shocks, both domestic and global, have played a major role, as have shocks to the exchange rate.

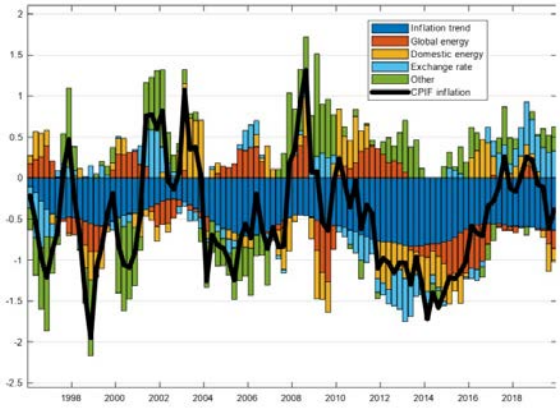


Figure 1: Historical decomposition of CPIF (the consumer price index with a fixed interest rate) inflation 1996Q1–2019Q4 using MAJA. Source: Corbo and Strid (2020).

A key purpose of MAJA and similar models used by central banks is to analyze the consequences of changes in the policy rate. The result of such analyses can be described in impulse-response diagrams. These track the developments over time of the endogenous variables after an unexpected change in the policy rate (or any other shock). Figure 2 below illustrates such responses, assuming that, after a policy shock occurs, the central bank reverts back to following their original rule. The shock is an unexpected hike in the nominal interest rate, i.e., a v_t shock (which gradually converges down to zero), and we see hump-shaped responses in real variables as well as in inflation: the initial response is modest and then builds, after which a peak is reached, then leading to a gradual return to the initial position. We note that the policy rate is also capable of changing the real exchange rate: an increase in the interest rate lowers the real exchange rate, i.e., it lowers the cost of foreign goods relative to domestic goods. Of course, even MAJA is stylized in many ways, so the counterfactual analyses it delivers only form part of

the inputs into how the policy rate is set in practice.

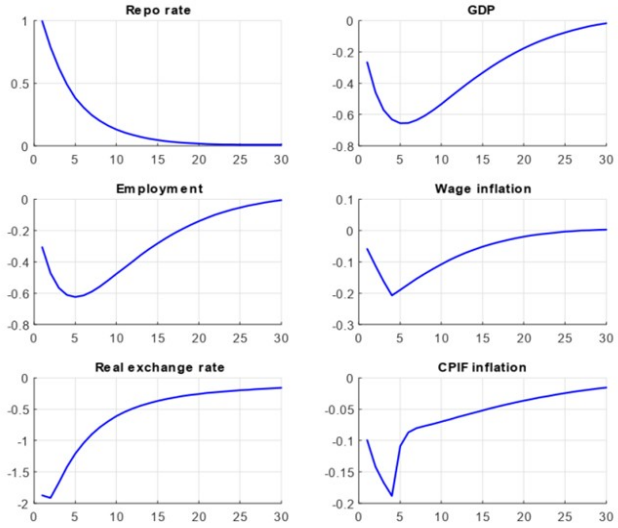


Figure 2: The effects of an increase in Sveriges Riksbank's policy rate.

Challenges The core New Keynesian model, as discussed above, along with many extended versions of it, is derived using linearization techniques. This is convenient, but it also risks missing important nonlinearities. An example of a nonlinearity is the zero lower bound, i.e., the fact that central banks view policy rates much below zero as infeasible. Even if this bound only applies in "non-normal" times, it will affect the dynamics of the model in normal times and this will not be captured if the model is solved with linearization methods. However, the methods for solving the model globally, taking account of such nonlinearities, have developed significantly over the past few decades.

The New Keynesian model builds on the rational expectations paradigm: all agents in the model are assumed to form expectations that are consistent with the predictions of the model.[57] The internal consistency provided by the requirement of rational expectations has proven extremely powerful and useful for modeling far outside the scope of New Keynesian models. However, a substantial problem for policymakers is the occasional occurrence of very hard-to-predict events—so-called black swans, such as a sudden war, a pandemic, or the election of an erratic politician. How such events change the expectations of agents in the real world is extremely hard, if not impossible, to capture within the present framework. Given a sudden burst

[57] It is important to note that this does not require the assumption that agents have full information; there are many extensions to the basic model where agents also face signal-extraction problems.

in, say, energy prices, the core model tends to predict immediate, albeit slow, convergence back toward the economy's original position, but we cannot be sure that consumers will not fear that inflation spirals out of control and hence act accordingly. Thus, the model presumes that inflation expectations are "anchored" around the target, but anchoring cannot be taken for granted. Of course, the models can still be, and very much have been in the context of recessions or sudden inflation bursts around the world, very useful for analyzing the transition from non-normal to normal times.[58]

The specific form of price or wage stickiness assumed in most New Keynesian models has also been criticized: these formulations, whether that by Calvo or Rotemberg, are ad hoc and not derived based on deeper primitives. This can be problematic especially in times of large, sudden increases in demand or in costs, as one would expect firms to then change prices more frequently and with more ease. Settings with arguably more plausible assumptions, such as a case where prices can always be changed at a fixed (possibly random), finite cost, have therefore been analyzed in the literature, but no new consensus model has yet emerged.[59]

Another weakness that has drawn significant attention is the so-called forward guidance puzzle. During the prolonged period of zero interest rates, Woodford and others argued that the central banks could communicate the interest-rate policy they planned to pursue later, including when the zero lower bound would not bind. The idea is that such communication could lower loan rates on longer horizons today and hence stimulate current demand, and the idea has been widely adopted. Del Negro et al. (2023) analyzed such an experiment using the simple New Keynesian model and found, surprisingly, that announcements far into the future had unreasonably strong effects on demand and prices today. Several possible solutions to the puzzle—in the form of amendments to the core model—have been suggested since then; arguably, these suggestions all have some merit, while no single explanation has emerged as a consensus.

5 Policy evaluation

The analysis of the source of business cycles, the way in which a variety of different shocks propagate through the economy, and how the monetary and fiscal authorities can and should stabilize the economy is clearly different today: central banks, fiscal agencies, and international organizations have adopted the New Keynesian approach to stabilization as a core tool. In con-

[58] Uribe (2022) argued that a version of the New Keynesian model with a randomly moving inflation target can account well for the data; here, de-anchoring can thus be given the interpretation of a changed target. Uribe also showed that a sudden increase in the nominal rate can lead inflation to rise quickly, should markets interpret the rate increase to mean a rise in the inflation target and hence accommodate it, and can even stimulate the economy, if the nominal rate rises slowly enough.

[59] For an early example, see Golosov and Lucas (2007).

crete terms, however, has the newer equipment, with microfounded models of the macroeconomy, helped us better stabilize business cycles and, if so, how?

Let us begin with "how." As our economy goes between its peaks and troughs, analysts—both academic researchers and practitioners—make a more concerted effort in identifying the microeconomic fundamentals behind the changes in our aggregates. There are many examples of this general point. One is how the slow recovery of employment after the Great Recession, which took place in the second part of the 2000s, was examined by looking at microeconomic data on labor market flows and the detailed workings of that market: had structural (technology or other) shocks made the matching process between workers and firms more difficult? Such inquiries often give important clues about appropriate policy reactions—one of which is laissez-faire—that can be analyzed using the new models. Another example is the recent rapid inflation burst in most of the Western world beginning in 2021: rather than merely studying the time series behavior of macroeconomic (perhaps especially monetary) aggregates, economists have studied the heterogeneity in price movements across space and sectors to understand what caused the shock, and what caused it to be somewhat persistent. Some of these studies also isolate new mechanisms that may lead us to enrich the New Keynesian framework further, at least for some applications, as well as to caution against mechanical use of the models, since they do have weaknesses.[60]

Another important area is the separate, but related, roles of monetary and fiscal policy in stabilization, where the microeconomic approach has made us focus more on the efficacy of different interventions. Fiscal policy aimed at increasing consumer demand may be very effective, and more so than monetary policy, if aimed at the right groups of consumers; monetary policy too may need to be directed for maximum impact, such as during the Covid-19 pandemic.

On a broader "how" level, the presumptions underlying the New Keynesian model include the existence of an independent central bank or, at least, that there is an ability to commit monetary policy to a rule. The analysis itself does not require such a stance; one could replace the interest-rate rule with a description of how politics may shape monetary policy decisions over time. At the same time, the use of an interest-rate rule with strong responses to inflationary pressures, along with responses to the output gap, has become an important benchmark for central banks. The New Keynesian framework points to important advantages of commitment and, relatedly, to transparency. In particular, Woodford (2005) argues forcefully for such a view: by making public its goals and methods, it increases its ability to conduct successful stabilization policy. Clearly, forward guidance is one such example: it is an "instrument" that

[60] As an example, a key insight we discuss above is that inflation expectations are very important and that we still do not have a full understanding of the circumstances under which their "anchoring" may become less firm and, if so, what appropriate measures to take.

fundamentally relies on forward-looking and a clear understanding of how the central bank operates.

Turning to the question of whether these models have been helpful, the evaluation is of course challenging for several reasons. One is that, since the New Keynesian model became available, the shocks our economies have experienced are not exactly the same—neither in their nature nor in their magnitude—as the shocks that occurred historically. Nevertheless, Eichengreen and O'Rourke (2010) have argued, based on aggregate data, that the rebound after the Great Recession was much faster than that after the Great Depression.[61] De Grauwe and Ji (2020) added the Covid-19 recession and argued similarly. One can also argue that the recent inflation shock that began in 2021, following increases in energy prices and supply-chain shocks in times of global unrest, is similar to the oil-price shock in 1973. Clearly, the rest of the 1970s—a prolonged period of time—were characterized by high inflation and low output growth (the "stagflation" period), whereas the recent inflation shock was not followed by persistently very high inflation, nor by stagnant output growth or persistently high unemployment. For examples of studies reaching this conclusion, based on data sets comprising many countries, see Ari et al. (2023), Forbes et al. (2024), Blanchard and Bernanke (2024), and Forbes et al. (2025).

One must also note that the New Keynesian framework is primarily about business-as-usual policy: how to stabilize given a set of exogenous shocks that have materialized. That is, it is not designed to prevent shocks from occurring in the first place. In the case of the Covid-19 episode, it is obvious that neither monetary nor fiscal policy could have prevented the shock. In contrast, the shock(s) that led to the Great Recession would, according to most analysts, at least to some extent have been possible to influence using advance changes in policy, though not through business-as-usual stabilization tools but rather through regulation.[62]

Nevertheless, the emphasis on forward-looking that is central to the New Keynesian paradigm has also allowed for new methods—such as transparency and forward guidance—that can make the propagation of shocks more muted and less persistent. Several studies, including Forbes et al. (2024), also emphasize central-bank independence as a relevant correlate in this context. In sum, while it is hard to prove beyond reasonable doubt that the relatively quick rebound after recessions in our economies over the past half century is due to better monetary and fiscal stabilization policy, there does seem to be a consensus that they have played a very important role.

[61] See also Almunia et al. (2010).

[62] Stricter regulation of financial markets can diminish the risk for excessive asset-price fluctuations, while of course also risking making these markets less effective on average.

6 Conclusions

This year's Prize is awarded to Robert E. Hall and Michael Woodford for their fundamental contributions in macroeconomics. Their contributions have generated new ideas and discoveries in a still vibrant research space, expanding the literature dramatically. The follow-up work builds directly on the original frameworks described in this document and can be seen as both robustness checks—the foundation is solid—and as incorporating new features that are interesting in themselves, with new, valuable insights as a result.

Together, Hall's and Woodford's contributions amount to a new foundation for the theory of macroeconomic fluctuations, allowing us to obtain a deeper understanding of the sources and propagation mechanisms behind recessions and booms. Using these new foundations, we are better positioned for analyzing and conducting stabilization policy. Initially, the focus of the research program on stabilization was monetary policy, but as a result of the new framework, it has also become the core tool for analyzing fiscal policy. The new framework is in widespread use—in research, in teaching, and by policymakers around the world—and is broadly viewed to have improved macroeconomic outcomes.

References

Agarwal, S. and W. Qian (2014). Consumption and debt response to unanticipated income shocks: Evidence from a natural experiment in Singapore. *American Economic Review 104*(12), 4205–4230.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics 109*(3), 659–684.

Akerlof, G. A. and J. L. Yellen (1985). A near-rational model of the business cycle, with wage and price inertia. *The Quarterly Journal of Economics 100*(Supplement), 823–838.

Almunia, M., A. Benetrix, B. Eichengreen, K. H. O'Rourke, and G. Rua (2010). From great depression to great credit crisis: similarities, differences and lessons. *Economic Policy 25*(62), 219–265.

Angeletos, G.-M., C. Lian, C. Wolf, and D. R. Zhang (2026). Monetary-fiscal interactions. Working paper.

Ari, A., C. Mulas-Granados, V. Mylonas, L. Ratnovski, and W. Zhao (2023). One hundred inflation shocks: Seven stylized facts. *IMF Working Paper 2023*(190).

Attanasio, O. P. and G. Weber (1993). Consumption growth, the interest rate and aggregation. *The Review of Economic Studies 60*(3), 631–649.

Azzimonti, M., P. Krusell, A. McKay, and T. Mukoyama (Eds.) (2026). *Macroeconomics*. New York, NY: Oxford University Press.

Barro, R. J. and D. B. Gordon (1983). Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics 12*(1), 101–121.

Benhabib, J., S. Schmitt-Grohé, and M. Uribe (2001a). Monetary policy and multiple equilibria. *American Economic Review 91*(1), 167–186.

Benhabib, J., S. Schmitt-Grohé, and M. Uribe (2001b). The perils of Taylor rules. *Journal of Economic Theory 96*(1-2), 40–69.

Blanchard, O. and J. Galí (2007). Real wage rigidities and the New Keynesian model. *Journal of Money, Credit and Banking 39*(s1), 35–65.

Blanchard, O. J. and B. S. Bernanke (2024, May). An analysis of pandemic-era inflation in 11 economies. Working Paper 32532, National Bureau of Economic Research.

Blanchard, O. J. and N. Kiyotaki (1987). Monopolistic competition and the effects of aggregate demand. *American Economic Review 77*(4), 647–666.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review 98*(5), 1887–1921.

Boehm, J., E. Fize, and X. Jaravel (2025). Five facts about MPCs: Evidence from a randomized experiment. *American Economic Review 115*(1), 1–42.

Browning, M. and M. D. Collado (2001). The response of expenditures to anticipated income changes: panel data estimates. *American Economic Review 91*(3), 681–692.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics 12*(3), 383–398.

Campbell, J. Y. and A. Deaton (1989). Why is consumption so smooth? *Review of Economic Studies 56*(3), 357–373.

Campbell, J. Y. and N. G. Mankiw (1989). Consumption, income, and interest rates: Reinterpreting the time series evidence. *NBER Macroeconomics Annual 4*, 185–216.

Campbell, J. Y. and N. G. Mankiw (1991). The response of consumption to income: A cross-country investigation. *European Economic Review 35*(4), 723–756.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy 113*(1), 1–45.

Clarida, R., J. Galí, and M. Gertler (1999). The science of monetary policy: a New Keynesian perspective. *Journal of Economic Literature 37*(4), 1661–1707.

Clarida, R., J. Galí, and M. Gertler (2001). Optimal monetary policy in open versus closed economies: an integrated approach. *American Economic Review 91*(2), 248–252.

Corbo, V. and I. Strid (2020). MAJA: A two-region DSGE model for Sweden and its main trading partners. Working Paper 391, Sveriges Riksbank, Stockholm.

Corsetti, G., L. Dedola, and S. Leduc (2008). International risk sharing and the transmission of productivity shocks. *The Review of Economic Studies 75*(2), 443–473.

Corsetti, G. and P. Pesenti (2001). Welfare and macroeconomic interdependence. *The Quarterly Journal of Economics 116*(2), 421–445.

Cuddington, J. T. (1982). Canadian evidence on the permanent income-rational expectations hypothesis. *The Canadian Journal of Economics/Revue canadienne d'Economique 15*(2), 331–335.

Daly, V. and G. Hadjimatheou (1981). Stochastic implications of the life cycle-permanent income hypothesis: Evidence for the UK economy. *Journal of Political Economy 89*(3), 596–599.

De Grauwe, P. and Y. Ji (2020). A tale of three depressions. VoxEU.org CEPS Policy Portal.

Del Negro, M., M. P. Giannoni, and C. Patterson (2023). The forward guidance puzzle. *Journal of Political Economy Macroeconomics 1*(1), 43–79.

Diamond, P. A. (1982). Aggregate demand management in search equilibrium. *Journal of Political Economy 90*(5), 881–894.

Dixit, A. K. and J. E. Stiglitz (1977). Monopolistic competition and optimum product diversity. *American Economic Review 67*(3), 297–308.

Eichengreen, B. and K. H. O'Rourke (2010, March). A tale of two depressions: What do the new data tell us? VoxEU.org.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance 25*(2), 383–417.

Fischer, S. (1977). Long-term contracts, rational expectations, and the optimal money supply rule. *Journal of Political Economy 85*(1), 191–205.

Flavin, M. A. (1981). The adjustment of consumption to changing expectations about future income. *Journal of Political Economy 89*(5), 974–1009.

Forbes, K., J. Ha, and M. A. Kose (2024). Rate cycles. Technical report, MIT Sloan Research Paper No. 7185-24.

Forbes, K., J. Ha, and M. A. Kose (2025, May). Tradeoffs over rate cycles: Activity, inflation and the price level. Working Paper 33825, National Bureau of Economic Research.

Friedman, M. (1957). *A Theory of the Consumption Function*. Princeton, NJ: Princeton University Press.

Galí, J. and T. Monacelli (2005). Monetary policy and exchange rate volatility in a small open economy. *The Review of Economic Studies 72*(3), 707–734.

Golosov, M. and R. E. Lucas, Jr. (2007). Menu costs and Phillips curves. *Journal of Political Economy 115*(2), 171–199.

Goodfriend, M. and R. G. King (1997). The new neoclassical synthesis and the role of monetary policy. *NBER Macroeconomics Annual 12*, 231–283.

Gray, J. A. (1976). Wage indexation: a macroeconomic approach. *Journal of Monetary Economics 2*(2), 221–235.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence. *Journal of Political Economy 86*(6), 971–987.

Hall, R. E. (1986). The role of consumption in economic fluctuations. In R. J. Gordon (Ed.), *The American business cycle: Continuity and change*, pp. 237–266. Chicago: University of Chicago Press.

Hall, R. E. (1988). Intertemporal substitution in consumption. *Journal of Political Economy 96*(2), 339–357.

Hall, R. E. and F. S. Mishkin (1982). The sensitivity of consumption to transitory income: Estimates from panel data on households. *Econometrica 50*(2), 461–481.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica 50*(4), 1029–1054.

Hansen, L. P. and K. J. Singleton (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica 50*(5), 1269–1286.

Hansen, L. P. and K. J. Singleton (1983). Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy 91*(2), 249–265.

Havranek, T. and A. Sokolova (2020). Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say "probably not". *Review of Economic Dynamics 35*, 97–122.

Hendren, N. (2017). Knowledge of future job loss and implications for unemployment insurance. *American Economic Review 107*(7), 1778–1823.

Hsieh, C.-T. (2003). Do consumers react to anticipated income changes? evidence from the Alaska Permanent Fund. *American Economic Review 93*(1), 397–405.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control 17*(5-6), 953–969.

Hymans, S. H. and H. T. Shapiro (1976). The allocation of household income to food consumption. *Journal of Econometrics 4*(2), 167–188.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review 96*(5), 1589–1610.

Johnson, P. (1983). Life-cycle consumption under rational expectations: Some Australian evidence. *Economic Record 59*(4), 345–350.

Kennan, J. (1979). The estimation of partial adjustment models with rational expectations. *Econometrica 47*(6), 1441–1455.

Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. London: Macmillan. 14th edition, 1973.

Kimball, M. S. (1995). The quantitative analytics of the basic neomonetarist model. *Journal of Money, Credit and Banking 27*(4), 1241–1277.

Kiyotaki, N. (1985). Macroeconomics of monopolistic competition. Harvard University Ph.D. dissertation.

Kueng, L. (2018). Excess sensitivity of high-income consumers. *The Quarterly Journal of Economics 133*(4), 1693–1751.

Kuznets, S. (1962). Quantitative aspects of the economic growth of nations: VII. The share and structure of consumption. *Economic Development and Cultural Change 10*(2, Part 2), 1–92.

Kydland, F. E. and E. C. Prescott (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy 85*(3), 473–491.

Kydland, F. E. and E. C. Prescott (1982). Time to build and aggregate fluctuations. *Econometrica 50*(6), 1345–1370.

Lucas, Jr., R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory 4*(2), 103–124.

Lucas, Jr., R. E. (1976). Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, Volume 1, pp. 19–46. North-Holland.

MaCurdy, T. E. (1983). A simple scheme for estimating an intertemporal model of labor supply and consumption in the presence of taxes and uncertainty. *International Economic Review 24*(2), 265–289.

Mankiw, N. G. (1981). The permanent income hypothesis and the real interest rate. *Economics Letters 7*(4), 307–311.

Mankiw, N. G. (1982). Hall's consumption hypothesis and durable goods. *Journal of Monetary Economics 10*(3), 417–425.

Mankiw, N. G. (1985a, 08). Consumer durables and the real interest rate. *The Review of Economics and Statistics 67*(3), 353–362.

Mankiw, N. G. (1985b). Small menu costs and large business cycles: A macroeconomic model of monopoly. *The Quarterly Journal of Economics 100*(2), 529–538.

McCallum, B. T. (1976a). Rational expectations and the estimation of econometric models: An alternative procedure. *International Economic Review 17*(2), 484–490.

McCallum, B. T. (1976b). Rational expectations and the natural rate hypothesis: some consistent estimates. *Econometrica 44*(1), 43–52.

Muellbauer, J. (1983). Surprises in the consumption function. *The Economic Journal 93*(Supplement), 34–50.

Muth, J. F. (1960). Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association 55*(290), 299–306.

Parker, J. A. (1999). The reaction of household consumption to predictable changes in social security taxes. *American Economic Review 89*(4), 959–973.

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review 103*(6), 2530–2553.

Paxson, C. H. (1993). Consumption and income seasonality in Thailand. *Journal of Political Economy 101*(1), 39–72.

Phelps, E. S. (1967). Phillips curves, expectations of inflation and optimal unemployment over time. *Economica 34*(135), 254–281.

Phelps, E. S. and J. B. Taylor (1977). Stabilizing powers of monetary policy under rational expectations. *Journal of Political Economy 85*(1), 163–190.

Phillips, A. W. (1958). The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861-1957. *Economica 25*(100), 283–299.

Rotemberg, J. J. (1982a). Monopolistic price adjustment and aggregate output. *The Review of Economic Studies 49*(4), 517–531.

Rotemberg, J. J. (1982b). Sticky prices in the United States. *Journal of Political Economy 90*(6), 1187–1211.

Rotemberg, J. J. (1987). The New Keynesian microfoundations. *NBER Macroeconomics Annual 2*, 69–104.

Rotemberg, J. J. and M. Woodford (1991). Markups and the business cycle. *NBER Macroeconomics Annual 6*, 63–129.

Rotemberg, J. J. and M. Woodford (1992). Oligopolistic pricing and the effects of aggregate demand on economic activity. *Journal of Political Economy 100*(6), 1153–1207.

Rotemberg, J. J. and M. Woodford (1996a). Imperfect competition and the effects of energy price increases on economic activity. *Journal of Money, Credit and Banking 28*(4), 549–577.

Rotemberg, J. J. and M. Woodford (1996b). Real-business-cycle models and the forecastable movements in output, hours, and consumption. *The American Economic Review 86*(1), 71–89.

Rotemberg, J. J. and M. Woodford (1997). An optimization-based econometric framework for the evaluation of monetary policy. *NBER Macroeconomics Annual 12*, 297–346.

Rotemberg, J. J. and M. Woodford (1999). The cyclical behavior of prices and costs. *Handbook of Macroeconomics 1*, 1051–1135.

Rubbo, E. (2023). Networks, Phillips curves, and monetary policy. *Econometrica 91*(4), 1417–1455.

Sargent, T. J. (1978). Rational expectations, econometric exogeneity, and consumption. *Journal of Political Economy 86*(4), 673–700.

Sargent, T. J. and N. Wallace (1975). "rational" expectations, the optimal monetary instrument, and the optimal money supply rule. *Journal of Political Economy 83*(2), 241–254.

Shapiro, M. D. (1984). The permanent income hypothesis and the real interest rate: some evidence from panel data. *Economics Letters 14*(1), 93–100.

Smets, F. and R. Wouters (2003). An estimated dynamic stochastic general equilibrium model of the euro area. *Journal of the European Economic Association 1*(5), 1123–1175.

Smets, F. and R. Wouters (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review 97*(3), 586–606.

Souleles, N. S. (1999). The response of household consumption to income tax refunds. *American Economic Review 89*(4), 947–958.

Svensson, L. E. O. (1986, 07). Sticky goods prices, flexible asset prices, monopolistic competition, and monetary policy. *The Review of Economic Studies 53*(3), 385–405.

Taylor, J. B. (1979). Staggered wage setting in a macro model. *The American Economic Review 69*(2), 108–113.

Taylor, J. B. (1980). Aggregate dynamics and staggered contracts. *Journal of Political Economy 88*(1), 1–23.

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy 39*, 195–214.

Uribe, M. (2022). The Neo-Fisher effect: Econometric evidence from empirical and optimizing models. *American Economic Journal: Macroeconomics 14*(3), 133–162.

Wicksell, K. (1898). *Geldzins und Güterpreise: Eine Studie über die den Tauschwert des Geldes bestimmenden Ursachen*. Jena: Gustav Fischer.

Wicksell, K. (1936). *Interest and Prices: A Study of the Causes Regulating the Value of Money*. London: Macmillan. Original work published 1898.

Woodford, M. (1994). Monetary policy and price level determinacy in a cash-in-advance economy. *Economic Theory 4*(3), 345–380.

Woodford, M. (1995). Price-level determinacy without control of a monetary aggregate. In *Carnegie-Rochester conference series on public policy*, Volume 43, pp. 1–46. Elsevier.

Woodford, M. (1996, July). Control of the public debt: A requirement for price stability? Working Paper 5684, National Bureau of Economic Research.

Woodford, M. (1998). Doing without money: controlling inflation in a post-monetary world. *Review of Economic Dynamics 1*(1), 173–219.

Woodford, M. (1999a). Commentary: How should monetary policy be conducted in an era of price stability? In *New Challenges for Monetary Policy*, pp. 277–316. Kansas City: Federal Reserve Bank of Kansas City.

Woodford, M. (1999b, July). Optimal monetary policy inertia. Working Paper 7261, National Bureau of Economic Research.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.

Woodford, M. (2005). Central bank communication and policy effectiveness. *Proceedings - Economic Policy Symposium - Jackson Hole* (Aug), 399–474.

Yogo, M. (2004). Estimating the elasticity of intertemporal substitution when instruments are weak. *Review of Economics and Statistics 86*(3), 797–810.

Yun, T. (1996). Nominal price rigidity, money supply endogeneity, and business cycles. *Journal of Monetary Economics 37*(2-3), 345–370.

Zeldes, S. P. (1989a). Consumption and liquidity constraints: an empirical investigation. *Journal of Political Economy 97*(2), 305–346.

Zeldes, S. P. (1989b). Optimal consumption with stochastic income: Deviations from certainty equivalence. *The Quarterly Journal of Economics 104*(2), 275–298.