Heterogeneous Experience and Constant-Gain Learning*

John Duffy† Michael Shin‡

February 8, 2024

Abstract

Recent evidence suggests that agents may base their forecasts for macroeconomic variables mainly on their personal life experiences. We connect this behavior to the concept of constant-gain learning (CGL) in macroeconomics. Our approach incorporates both heterogeneity in the life cycle via the perpetual youth model and learning from experience (LfE) into a linear expectations model where agents are born and die with some probability every period. For LfE, agents employ a decreasing-gain learning (DGL) model using data only from their own lifetimes. While agents are using DGL individually, we show that in the aggregate, expectations follow an approach related to CGL, where the gain is now tied to the probabilities of birth and death. We provide a precise characterization of the relationship between CGL and our model of perpetual youth learning (PYL) and show that PYL can well approximate CGL while pinning down the gain parameter with demographic data. Calibrating the model to U.S. demographics leads to gain parameters similar to those found in the literature. Further, variation in birth and death rates across countries and time periods can help explain the empirical time-variation in gains. Finally, we show that our approach is robust to alternative ways of modeling individual agent learning.

Keywords: Bounded rationality, Learning, Experience, Heterogeneity, Perpetual-youth model

JEL Codes: D83, D84, E71

*We thank the editor and three anonymous referees for helpful comments on an earlier draft. We also thank Albert Marcet, Alex Grimaud, Bill Branch, Chris Gibbs, Fabio Milani, Nawid Siassi, and participants at the BSE Workshop in Expectations in Dynamic Macroeconomic Models and the 5th Bamberg Behavioral Macroeconomics Workshop for their valuable feedback.

†Department of Economics, University of California, Irvine and ISER, Osaka University. Email: duffy@uci.edu

‡Discipline of Finance, The University of Sydney. Email: michael.shin@sydney.edu.au

1 Introduction

Recent evidence suggests that households and firms form expectations in a manner that deviates from full information rational expectations (FIRE) (Malmendier and Nagel (2011, 2016), Coibion and Gorodnichenko (2015)). Specifically the evidence suggests that households overweight data from their own lifetimes and underweight data from periods prior to their birth. This is sometimes referred to as the "experience hypothesis", with extensive psychological underpinnings (Weber et al. (1993), Hertwig et al. (2004)).

Moreover, Malmendier and Nagel (2016) suggest that "*learning from experience (LfE) behaves approximately in the data as if it were generated by a constant-gain learning algorithm*" (Malmendier and Nagel, 2016, p. 79), and they provide some qualitative evidence in support of this claim. Our aim in this paper is to take a deeper dive into this argument and make a more explicit connection between the learning from "lived" experience approach to expectation formation and the use of constant-gain learning (CGL) algorithms in the literature on learning in macroeconomics (Evans and Honkapohja (2001), Cho et al. (2002) and Williams (2019)). In particular, we show precisely the degree to which learning from experience (LfE) and constant-gain learning (CGL) are substitutes for one another.

The learning literature in macroeconomics has sought to replace FIRE with the idea that agents are boundedly rational. For instance, Evans and Honkapohja (2009) proposes the "cognitive consistency principle", which requires that agents in macroeconomic models should not be more knowledgeable than (good) economists. In practice, this means that agents run least squares regressions of the variables they are seeking to forecast on *all* past data to form forecasts of those variables. But good "econometric learning" using the *entire* history of past data might not be what households are doing.

In this paper, we imagine that households use decreasing-gain learning (DGL) algorithms as in the literature on least squares learning, but they limit the data they use to their *own personal histories*, which vary with the date of their birth, following the evidence of Malmendier and Nagel (2011). In effect, they discount (with factor 0) all data that occurred before they were born (or became adults). In its recursive formulation, least squares learning involves a decreasing gain on the most recent error term, so this approach is known as the decreasing-gain learning (DGL) approach. An advantage of DGL is that if a rational expectations equilibrium (REE) is expectationally stable (E-stable) and the agents' perceived law of motion (PLM) includes the REE, then DGL will converge, in the limit to that REE (Marcet and Sargent (1989), Evans and Honkapohja (2001)). An alternative approach, which often yields a good fit to macroeconomic data, is the constant-gain learning (CGL) approach where the gain on the most recent error term is some exogenously chosen constant. As a consequence, CGL systems remain volatile and ever-vigilant to external shocks or structural breaks in the data-generating process (DGP). This makes CGL quicker to adjust to such shocks in contrast to DGL systems, which are slower to respond to changes.

While CGL has been a benchmark model in the learning literature, it has sometimes been criticized as being *ad-hoc*, with no clear correspondence of the gain parameter with the data (Adam and Marcet (2011)).[1] Here, we take these critiques seriously and we propose an alternative approach. We ask whether a heterogeneous agent, econometric learning model where individual agents are learning from experience (LfE) using data only from their own lifetimes, can aggregate up to and approximate the CGL approach. In effect, we aim to provide a microfounded, demographic-based rationalization for the use of CGL.

We illustrate our idea using a simple linear expectations model to formulate the rational expectations (RE) hypothesis. For further tractability, we assume that agents learn using a DGL method related to least squares learning known as stochastic gradient learning (Evans and Honkapohja (2001)). Agents form expectations, within a Blanchard-Yaari perpetual youth model, where old agents are born and die with some random probabilities each period. To incorporate learning from experience (LfE), we assume that agents only weigh data observed within their own lifetimes.[2]

We provide an analytical characterization for the learning rule under our "perpetual youth, learning from experience" (PYL) model. Under our approach, the constant-gain parameter of the aggregated learning system becomes a function of the demographics of the model, specifically the birth and death rates. We show that there is a difference between the PYL model weights and those of the standard CGL framework without learning from experience (LfE). In particular, while CGL leads to geometrically declining weights on past data, our perpetual youth, learning from experience model leads to a mathematical formulation known as a Lerch transcendent on past data (Abramowitz and Stegun (1964)). Specifically, we show that PYL amounts to a functional transformation of the CGL weights on past data multiplied by the Lerch transcendent.

A key contribution of our PYL model is that it removes the free parameter from the CGL model by tying the gain parameter to the empirical demographics of the population. As noted, PYL depends on the *birth and death rates* in the Blanchard-Yaari model with

[1] Various modifications to constant-gain learning (CGL) have been proposed such as Marcet and Nicolini (2003) and Milani (2014). More recently, Berardi (2020) proposes a Bayesian interpretation of the gain parameter as the probability that there are changes in the estimates of the learning variables.

[2] To clarify, agents will also incorporate the most recent observation as well. With proper timing protocols, the most recent observation x_{t-1} can be considered to be *within* an agent's lifetime.

lower birth and death rates leading to more weight on data far into the past. We also show that PYL nests two well-known learning models: DGL and naïve learning as special cases. In the stationary, no population growth case, as the birth and death rates go to 0 so that eventually no agent dies in the model, PYL converges to DGL. As the birth and death rates go to 1, so that agents only live for 1 period, PYL converges to naïve expectations.

As there is no objective way to compare the CGL and PYL models since CGL relies on a constant-gain parameter, while PYL is pinned down by the birth and death rates, we initialize both models by setting the first period's expectations under PYL equal to those under CGL. We then run simulations of our PYL model using a standard linear asset pricing model and compare the moments from those simulations with the CGL model. We show that, while there is no exact mapping between the PYL and CGL learning models, the PYL model approximates CGL in that the simulated moments are approximately the same for both models. In particular, while we find some differences in the weights between the two models that we can explicitly characterize, in simulations we find essentially no differences in the first moments and small differences in higher-order moments under reasonable parameterizations.

We also calibrate our PYL model to demographic data and find that under current U.S. demographics, the implied constant gain is 0.018. This value is consistent with estimates reported in some key papers in the empirical macroeconomic learning literature where the constant gain is found to range between 0.01 - 0.04 (Orphanides and Williams (2005), Milani (2007), Eusepi and Preston (2018)). We also compute the implied gains across different countries and time periods and find that variations in birth and death rates lead to an interesting cross-section of gains. In particular, countries with low birth rates have lower gains than countries with higher birth rates. Hence, our model also predicts both cross-country and time-period differences in implied learning gain parameters. We suggest that our model can be used to understand findings in the literature such as differences in the anchoring of inflation expectations across countries (Coibion and Gorodnichenko (2015)) and cross-country variations in gains (Branch and Evans (2006), Slobodyan and Wouters (2012)).

Finally, we consider how our results would change if we replaced the assumption that individual agents learn using a DGL approach, with three other variants for individual learning: (1) adding additional data, (2) a generalized DGL, and (3) biased weights (Malmendier and Nagel (2016)). We show that with these learning rules our PYL approach can still deliver effective constant gains that are within the range of values found in the literature.

2 Literature Review

This paper contributes to several literatures. First, it contributes to the literature on adaptive learning in macroeconomics. In particular, we contribute to both the literature on alternative gain schemes to DGL and CGL in learning and to the literature on the calibration and estimation of gain parameters. Marcet and Sargent (1989) and Evans and Honkapohja (2001) provide the first systematic relaxations of the RE hypothesis and replace it with an econometric learning approach. Marcet and Nicolini (2003) provide a learning model where agents switch to constant-gain learning (CGL) if the forecast errors of the agent's learning model are higher than some threshold. Sargent and Williams (2005) show that a Kalman filter can be well approximated with a CGL model under certain priors. Evans et al. (2010) show that a model with stochastic gradient learning and a constant gain can be optimal when agents suspect parameter drift.

A closely related paper is, Adam and Marcet (2011), who relax the rational expectations hypothesis (REH), where agents are "internally rational", i.e., they maximize expected utility given consistent future subjective beliefs updated in a Bayesian way, but differently from REH, these agents may not be "externally rational", i.e., knowing the true data-generating process (DGP). They use internal rationality in an asset pricing model and show that the optimality conditions only involve today's price and the expected future price and dividend in the next period via a one-step-ahead pricing equation where expectations are replaced by a least square learning algorithm if prior beliefs are close to the RE beliefs (Adam et al. (2016), Adam et al. (2017)). Our approach also employs one-step-ahead expectations formation, but we consider the aggregate implications of such updating for agents with heterogeneous data histories.

We also contribute to the literature on empirical estimates of the constant-gain parameter for macroeconomic learning systems. The key papers in this literature are Orphanides and Williams (2005), Branch and Evans (2006), Milani (2007, 2008), Eusepi and Preston (2011), Slobodyan and Wouters (2012), and Berardi and Galimberti (2017). We differ from this literature in that we do not estimate the gain parameter but instead, we provide a method for pinning down the implied constant-gain parameter that arises from our PYL model using empirical birth and death rates.

Our paper also contributes to the literature on experience-based learning in that we relate the empirical findings from that literature to constant-gain learning (CGL). The experience-based learning approach begins with a key paper by Malmendier and Nagel (2011) who proposes that macroeconomic events over an individual's lifetimes disproportionately affect their investment choices. (See Malmendier (2021) for a survey). Their key means of identifying this "experience hypothesis" is to relate differences in risk attitudes between old and young people with differences in their lifetime experiences. Individuals who experience periods of low stock-market returns should express a lower willingness to participate in the stock market and invest less in stocks and they verify this using measures of risk aversion and household portfolio data. Our paper takes the experience hypothesis seriously and utilizes a stylized version where agents assign zero weight to data that occurred before they were born.

While Malmendier and Nagel speculate that a constant-gain learning (CGL) model could be a good approximation to experience-based learning, their comparison is mainly qualitative. Our approach provides a different micro-foundation for experience-based learning as we employ the perpetual youth model to allow for belief and experience heterogeneity, which makes the model more tractable. Following that approach, we provide an explicit analytic characterization of the relationship between constant-gain learning (CGL) and our perpetual youth, learning from experience (LfE) model. Shin (2021) shows that learning from experience (LfE) can be replicated in the laboratory in an asset pricing experiment with negative experiences being quantitatively stronger than positive ones. We note that Malmendier et al. (2020) use a multi-period overlapping generations model with LfE to explain their earlier empirical findings through an asset pricing model.

Finally, the paper most closely related to this one is by Nakov and Nuño (2015). The paper follows Adam and Marcet (2011) and incorporate "internal rationality" in an asset pricing model but they adopt a heterogeneous agent, "learning from experience" approach following the work of Malmendier and Nagel (2016) in which agents learn only from their own experience. They use a Blanchard-Yaari model where individuals learn about both prices and dividends and update beliefs in a Bayesian fashion. They show their model can generate stock prices that exhibit patterns similar to those in the data and that the dynamics of average beliefs in their model can be approximated by a constant-gain learning scheme where the average gain is a nonlinear function of the survival rate and the individual gain parameter. Our paper is complementary to theirs but differs in key respects. First, we provide an exact characterization of the relationship between demographic factors and the constant-gain parameter and not an approximation. Second, we show that both birth and death rates i.e. the population dynamics play an important role, whereas Nakov and Nuño (2015) consider only a single survival rate. Third, our main methodology is also used to connect alternative learning rules to a constant-gain parameter.

3 The Model

We use a simple linear expectations model as our laboratory.[3] The main equation takes the form:

$$x_t = \alpha + \beta x_t^e + \delta' w_{t-1} + \epsilon_t, \tag{1}$$

where x_t is the endogenous variable at time t, α is a constant term, $\beta \in [0, 1)$ is the coefficient term on beliefs, x_t^e denotes aggregate expectations for the endogenous variable at time t, formed at the beginning of time t, δ' is a 1 by n vector of coefficients, w_{t-1} is an n by 1 vector of exogenous variables, and $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ is a noise term. This is a temporary equilibrium formulation in that expectations here can be taken as given and are not necessarily equilibrium objects as in a rational expectations equilibria (REE). The model is closed under a specification for expectations formation. Under rational expectations (RE), the unique REE solution is given by:

$$x_t = \overline{a} + \overline{b}' w_{t-1} + \epsilon_t \quad \text{where} \quad \overline{a} = (1-\beta)^{-1}\alpha \quad \text{and} \quad \overline{b} = (1-\beta)^{-1}\delta. \tag{2}$$

3.1 Blanchard-Yaari Distribution

Time is discrete. There is a continuum of agents $k \in [0, 1]$ indexed by both age i, and time, t. Specifically, agent $k_{i,t}$ is of age i at time t. Agents are born at age 0, and as will be clarified in the next section, advance one age at the start of each new period. Thus a new agent forms their first expectation $x_{i,t}^e$, at the beginning of time t, at age 1.

As in the Blanchard-Yaari model, agents are born at the birth rate $g \geq 0$ and die at the death rate $q \in (0, 1)$. We also assume that agents take an anticipated utility approach (Kreps (1998)) as in much of the learning literature where individuals use their *subjective* beliefs to estimate the likelihood of different outcomes and update these beliefs as new data arises.

We adopt a simple setup to make as clear as possible the differences and similarities between CGL and our perpetual youth learning approach. The perpetual youth environment creates a distribution of agents with different information sets. The number of age i agents at time $t > 0$ for $i \leq t$ is given by:

[3]This model can be derived from Muth's cobweb model as in Bray and Savin (1986) or from other standard models such as a linear asset pricing model or from a simple New Keynesian (NK) model.

$$N_{i,t} = g(1+g-q)^{t-i}(1-q)^{i-1}$$

with i-year olds for any $i \in \{1, 2, ..., t-1\}$. We can get the proportion of age i agents, $n_{i,t}$ by normalizing the population $N_{i,t}$ by the growth factor $(1+g-q)^t$ which leads to:[4]

$$n_{i,t} = \frac{g(1-q)^{i-1}}{(1+g-q)^i} \tag{3}$$

Note that Equation (3) is a proper distribution and integrates to 1. Here, since each agent of a given age i is homogeneous within their age group i, to save on notation, we can omit the agent's index k, as i will index both the individual agent's beliefs and their age group's beliefs. We are also free to drop the time subscript as $n_{i,t}$ is no longer a function of t. Average life expectancy in the model is given by $\frac{1}{q}$.

3.2 Timing Protocol

Each time period t is separated into three sub-periods. In sub-period 1, all agents age 1 year and form their expectations about the endogenous variable, $x^e_{i,t}$. In sub-period 2, markets clear and the endogenous variable, x_t, is realized. In sub-period 3, some agents die while others are born. Because sub-period 3 is still a component of time period t, we suppose that the time t value for the endogenous variable, x_t is included in the information set of agents born in period t.

3.3 Perpetual Youth Learning (PYL)

We assume that agents form imperfect expectations, which is formalized by allowing an agent's information set to only include x_t from their own lifetimes. Specifically, an agent of age i will have the following information set: $\{x_j\}_{j=t-i}^{t-1}$. For example, an agent of age 3 will have: $\{x_{t-1}, x_{t-2}, x_{t-3}\}$. When forming expectations for $x^e_{i,t}$, we assume that agents use a decreasing-gain learning (DGL) approach, that is, they place equal weight on all observations *within* their own lifetimes. DGL learning is a standard assumption in the learning literature as it has a close correspondence with ordinary least squares (OLS) regression.

We further assume that all agents have a correctly specified perceived law of motion (PLM) that nests the REE, Equation (2) as a special case. Agent i's PLM as to how the

[4]Note that as $t \to \infty$, $N_t \to \infty$ if $g > q$, hence, Equation (3) is valid as long as we are not at the limit, as the population will just be arbitrarily large (or small if $g < q$). In the case where $g = q$, Equation (3) holds even in the limit.

economy evolves is given by:

$$x_{i,t}^{e} = a_{i,t-1} + b' w_{t-1} \tag{4}$$

where $b' = \bar{b}$ is a 1 by n coefficient vector that is assumed to be perfectly known. We further assume that the entire history of the exogenous variables w_{t-1} is also known and that individuals do not limit their information sets only to realizations of w_{t-1} observed within their lifetime. The latter assumptions are made for tractability purposes.[5] We provide a formulation with learning about the coefficients on the exogenous variables w_{t-1}, in the Appendix and specify the analytical difficulties.

Thus, we focus on learning about the constant term, $a_{i,t}$, using only the history of x_t observed over each agent's lifetime as this makes our analytic results most clear. We use stochastic gradient learning as the DGL model instead of least squares learning for further tractability, although these learning rules are equivalent when b is known.[6] More specifically, we suppose that agents update the parameter $a_{i,t}$ using the stochastic gradient learning approach:

$$a_{i,t} = a_{i,t-1} + \frac{1}{t(i)} \left[x_t - a_{i,t-1} - b' w_{t-1} \right] \tag{5}$$

Since an agent's information set depends only on their age, we can replace $\frac{1}{t(i)}$ in the DGL weight with $\frac{1}{i}$. We can index these expectations by age as follows:

$$\begin{aligned}
a_{1,t} &= a_{1,t-1} + \frac{1}{1}[x_t - a_{1,t-1} - b' w_{t-1}] \ \text{(Age 1)} \\
a_{2,t} &= a_{2,t-1} + \frac{1}{2}[x_t - a_{2,t-1} - b' w_{t-1}] \ \text{(Age 2)} \\
a_{3,t} &= a_{3,t-1} + \frac{1}{3}[x_t - a_{3,t-1} - b' w_{t-1}] \ \text{(Age 3)} \\
&\cdots \\
a_{i,t} &= a_{i,t-1} + \frac{1}{i}[x_t - a_{i,t-1} - b' w_{t-1}] \ \text{(Age i)}
\end{aligned}$$

[5] Agents could also limit their history of w_{t-1} via learning from experience (LfE) but this has no significant impact on the analytical results.

[6] See Evans et al. (2010). Unlike least squares learning, the stochastic gradient learning algorithm does not require information on the second moments of the data, using an identity indexing matrix instead.

3.4 Perpetual Youth Learning Weights

Here, we solve for aggregate expectations, x_t^e. For PYL, let the population of agents of age i be denoted by $n_i = \frac{g(1-q)^{i-1}}{(1+g-q)^i}$. Then, we can define aggregate expectations by:

$$x_t^e \equiv \sum_{i=1}^{\infty} n_i x_{i,t}^e. \tag{6}$$

In order to solve for aggregate expectations, we need to aggregate across each individual's expectations.

3.4.1 Solving for Aggregate Expectations

There are two ways of solving for aggregate expectations: (1) rearranging infinite sums and (2) using matrices. Both are equivalent. The approach of rearranging terms in the infinite sums is simpler, and so we present that approach here. However, the use of matrices provides a more visual intuition while also showing a connection between our main results and to models where agents use the full history of data. Therefore, in the Online Appendix, we provide the derivation using matrices and the case with no regressors as this can provide stronger intuition.

We use the insight that the DGL rule under stochastic gradient learning is essentially a weighted average of past observations. The main observation here is that with known constant b', the forecasting task of agents is to just estimate a sample mean of past observations relative to their own history. It turns out the variable of interest is just a "residualized mean". That is, agents try to estimate the mean by taking out the regressor term i.e., their estimate is given by $x_t - b'w_{t-1}$. We call this term the "residualized mean".

More specifically,

$$\begin{aligned} a_{i,t} &= a_{i,t-1} + \frac{1}{i}[x_t - a_{i,t-1} - b'w_{t-1}] \\ &= \frac{1}{i}[(x_t - b'w_{t-1}) + (x_{t-1} - b'w_{t-2}) + ... + (x_{t-i+1} - b'w_{t-i})] \end{aligned}$$

where the equivalence follows from the fact that the recursive least squares (RLS) algorithm in Equation (5) can be rewritten as the sample mean of past observations by definition.

Plugging this into the PLM in Equation (4) and the resulting expression into (6) we get:

$$\begin{aligned} x_t^e &= \sum_{i=1}^{\infty} n_i(a_{i,t-1} + b'w_{t-1}) \\ &= b'w_{t-1} + \sum_{i=1}^{\infty} n_i a_{i,t-1} \\ &= b'w_{t-1} + \sum_{i=1}^{\infty} n_i \frac{1}{i}[(x_{t-1} - b'w_{t-2}) + (x_{t-2} - b'w_{t-3}) + ... + (x_{t-i} - b'w_{t-i-1})] \end{aligned}$$

where the second line above follows from the fact that the population sums to 1 and the b' terms are the same for every agent. Then we can define the second part of the equation as a_{t-1}:

$$a_{t-1} = \sum_{i=1}^{\infty} n_i \frac{1}{i}[(x_{t-1} - b'w_{t-2}) + (x_{t-2} - b'w_{t-3}) + ... + (x_{t-i} - b'w_{t-i-1})]$$

Here there are essentially two different infinite sums being multiplied together, the geometric-harmonic series $\sum \frac{n_i}{i}$ and the sum of past residualized means. Hence, a solution to this formulation is the same as finding the coefficients on the past $x_t - b'w_{t-1}$ terms by grouping like terms. It is well known that the geometric-harmonic series is convergent, so to be able to rearrange the terms, we need to show that the series is absolutely convergent. In particular, with sufficient regularity conditions on the residualized mean, we can guarantee absolute convergence, with one sufficient condition being that both x_t and w_t are stationary variables.[7]

Then, writing out the individual terms of the sum:

$$\begin{aligned} a_{t-1} = n_1(x_{t-1} - b'w_{t-2}) + \frac{1}{2}n_2\left[(x_{t-1} - b'w_{t-2}) + (x_{t-2} - b'w_{t-3})\right] + ... \\ + \frac{1}{i}n_i\left[(x_{t-1} - b'w_{t-2}) + (x_{t-2} - b'w_{t-3}) + ... + (x_{t-i} - b'w_{t-i-1})\right] + ... \end{aligned}$$

Note that every cohort places weight on the most recent observation $x_{t-1} - b'w_{t-2}$. Then the pattern is, every cohort besides the cohort of age 1 places weight on the observation $x_{t-2} - b'w_{t-3}$ and so on.

[7] In particular, the key assumption that is needed is that the discounted sum of future residualized means i.e., x_t and w_t are convergent. Another sufficient condition is that x_t and w_t have a growth rate that is bounded by some constant, c i.e., they are not growing faster than the denominator.

Rearranging, the terms, we get:

$$a_{t-1} = \left(\sum_{i=1}^{\infty} \frac{n_i}{i}\right)(x_{t-1} - b'w_{t-2}) + \left(\sum_{i=2}^{\infty} \frac{n_i}{i}\right)(x_{t-2} - b'w_{t-3}) + ... + \left(\sum_{j=i}^{\infty} \frac{n_j}{j}\right)(x_{t-j} - b'w_{t-j-1}) + ...$$

where we have to add an index j due to needing a second index number. In particular, we can define the solution for aggregate expectations x_t^e for the constant a_{t-1} in the aggregate perceived law of motion (PLM) as:

$$\begin{aligned} a_{t-1} &= S_1(x_{t-1} - b'w_{t-2}) + S_2(x_{t-2} - b'w_{t-3}) + ... \\ &= \sum_{i=1}^{\infty} S_i(x_{t-i} - b'w_{t-i-1}) \end{aligned} \tag{7}$$

where S_i is:

$$S_i = \sum_{j=i}^{\infty} \frac{n_j}{j} = \sum_{j=i}^{\infty} \frac{q(1-q)^{j-1}}{(1+g-q)^j} * \frac{1}{j} \tag{8}$$

Here we find that under PYL with learning from experience (LfE) with individual DGL, the aggregate weights S_i on past data are infinite sums of a geometric-harmonic series, starting from the element i instead of 1. Note that these weights are generally *not* geometrically declining as in CGL.[8]

3.5 Lerch Transcendent

We have an expression for aggregate expectations via Equations (7) and (8). To make further progress with analytic results we study a special case of the population distribution, Equation (3), where we set the birth rate equal to the death rate, $g = q$:

$$n_i^s = q(1-q)^{i-1} \tag{9}$$

Equation (9) is the stationary proportion of cohorts of age i with the superscript s for n^s to denote a stationary population. Note that in a stationary population, we can refer to either the birth rate g or the death rate q as they are both equal but we will refer to the death rate q, which again implies an average life expectancy of $\frac{1}{q}$ in the model, to help foster intuition. We will now characterize the solution using known analytical forms in the mathematics literature. Then we can characterize the S_i as taking elements of the partial

[8]It can be proven that the sums of geometric series are themselves, not geometric series (unless they are the same series) i.e., the class of geometric series are not closed under addition.

sums of the infinite sum for Equation (8):

$$
\begin{aligned}
S_i &= \sum_{j=i}^{\infty} \frac{q(1-q)^{j-1}}{j} \\
&= q(1-q)^{i-1} * \left(\sum_{k=0}^{\infty} \frac{(1-q)^k}{i+k} \right) \\
&= q(1-q)^{i-1} \Phi(1-q, 1, i)
\end{aligned}
$$

where the second equality above follows from splitting the series, factoring out $q(1-q)^{i-1}$, and setting the new index from k. Here, S_i is the ith partial sum of the geometric-harmonic series starting from $j = i$ and $\Phi(1-q, 1, i)$ is known as the Lerch transcendent (Abramowitz and Stegun (1964)), which is a generalization of the Riemann zeta function.[9]

The Lerch transcendent is the following function:

$$\Phi(z, s, a) \equiv \sum_{k=0}^{\infty} \frac{z^k}{(a+k)^s}$$

In essence, the Lerch transcendent is just an analytical generalization of the geometric-harmonic infinite series. For our example $z = 1-q, s = 1, a = i$:

$$\Phi(1-q, 1, i) = \sum_{k=0}^{\infty} \frac{(1-q)^k}{i+k}$$

Rewriting for expectations x_t^e under the partial sums S_i via Equation (7), we get:

$$
\begin{aligned}
x_t^e &= b'w_{t-1} + S_1(x_{t-1} - b'w_{t-2}) + S_2(x_{t-2} - b'w_{t-3}) + S_3(x_{t-3} - b'w_{t-4}) + ... \\
&= \sum_{i=1}^{\infty} S_i(x_{t-i} - b'w_{t-i-1}) + b'w_{t-1} \\
&= \sum_{i=1}^{\infty} (x_{t-i} - b'w_{t-i-1}) \left[q(1-q)^{i-1} \Phi(1-q, 1, i) \right] + b'w_{t-1} \\
&= \underbrace{\sum_{i=1}^{\infty} q(1-q)^{i-1}}_{\text{CGL Term}} \underbrace{\Phi(1-q, 1, i)}_{\text{PYL Term}} (x_{t-i} - b'w_{t-i-1}) + b'w_{t-1} \qquad (10)
\end{aligned}
$$

Thus, under PYL with $g = q$, aggregate expectations can be decomposed into a CGL term and a PYL term as in Equation (10).

[9] Note that $\Phi(1, s, 1) = \zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}$ is the famous Riemann zeta function.

In essence, aggregate expectations under PYL is a functional transformation of the CGL weights on past realizations multiplied by a Lerch transcendent. Again, as the CGL terms here refer to the population distribution, an alternative interpretation is that the PYL terms "twist" the CGL weights in an ordinary CGL model such that more weight is put on the older cohorts' beliefs.

3.6 Comparative Statics

In this section, we examine how the PYL weights, S_i, vary with the death rate q. The PYL weights are those that agents place on observation $t-i$ via Equation (10):

$$S_i = \sum_{j=i}^{\infty} q(1-q)^{j-1} * \frac{1}{j} = q(1-q)^{i-1}\Phi(1-q,1,i)$$

Of particular focus will be S_1 because of its role in the approximation in subsequent sections. Note, that we can also characterize S_i by iterating on the Lerch Transcendent formula. Further S_i terms will not have a simple algebraic form like with S_1. We can characterize S_1 with the following proposition:

Proposition 1. *The infinite sum of the following geometric-harmonic series is:*

$$S_1 = \sum_{i=1}^{\infty} \frac{g(1-q)^{i-1}}{(1+g-q)^i} * \frac{1}{i} = \frac{-g\ln(\frac{g}{1+g-q})}{1-q}. \tag{11}$$

In the case without population growth where $g=q$, we have:

$$S_1 = \sum_{i=1}^{\infty} \frac{q(1-q)^{i-1}}{i} = \frac{-q\ln q}{1-q} \tag{12}$$

The proof of Proposition 1 is in the Appendix.
The derivative of S_i with respect to the death rate q is:

$$\frac{\partial S_i}{\partial q} = (1-q)^{i-2}[\Phi(1-q,1,i)-1]$$

By Proposition 1, $S_1 = \frac{-q\ln q}{1-q}$. Then, the derivative of S_1 with respect to the death rate q is:

$$\frac{\partial S_1}{\partial q} = \frac{q-\ln q-1}{(1-q)^2} > 0, \;\; q \in (0,1)$$

Note that the sign of the derivative with respect to q of the general path of weights S_i depends on model parameters. That is because as the weights S_i shift to the left (as there are fewer old agents the death rate increases), some values of S_i will increase while others will decrease. More informative for our purposes will be to sign S_1, that is, what happens in general to the weight on earlier observations as the death rate q increases, which is positive for $q \in (0, 1)$. Intuitively, as q increases, the death rate goes up and there are more young and fewer old in the model. This increases the weight on recent observations which is an increase in S_1.

3.7 Connections to DGL and Naïve Learning

We also find that PYL with no population growth encompasses two well-known learning models as special cases: DGL and naïve learning. The learning rule for PYL is again as follows:

$$x_t^e = \sum_{i=1}^{\infty} q(1-q)^{i-1}\Phi(1-q, 1, i)(x_{t-i} - b'w_{t-i-1}) + b'w_{t-1}$$

Recall $q \in (0, 1)$ is the death rate here, which is the probability of death within the perpetual youth model, with $g = q$.

For the DGL case, when $q \to 0$, then no one in the model dies, which corresponds to the infinite horizon, representative agent case. DGL is the benchmark model in the learning literature as it is a natural analog to what econometricians do in practice. In that case, agents eventually weigh the entire history with $\frac{1}{t}$ as in the DGL model. We can think of this version of the model as shutting down the perpetual youth channel. For the naïve learning case, note that when $q \to 1$, the death rate is 100% and agents only live for 1 period. Hence, we only need only look at the first term weight in the equation above which is: $S_1 = q * \Phi(1-q, 1, 1) = \frac{-q \ln q}{1-q}$. Taking the limit as q approaches 1 gives:

$$\begin{aligned} x_t^e &= \lim_{q \to 1} q * \Phi(1-q, 1, 1)(x_{t-1} - b'w_{t-2}) + b'w_{t-1} \\ &= \lim_{q \to 1} \left(\frac{-q \ln q}{1-q} \right) (x_{t-1} - b'w_{t-2}) + b'w_{t-1} \\ &= (x_{t-1} - b'w_{t-2}) + b'w_{t-1} \end{aligned}$$

as $\lim_{q \to 1} \left(\frac{-q \ln q}{1-q} \right) = 1$ via L'hôpital's rule. Hence, the limit puts all the weight on the first term which amounts to naïve learning. Because the CGL model also converges to naïve learning as the gain parameter approaches 1, both models converge to naïve learning and are limiting cases of one another.

4 Comparison of PYL and CGL Models

In this section, we relate our PYL model to CGL as we wish to use our model to provide an approximation for CGL. We first note that the PYL formulation demonstrates there is no exact mathematical representation that directly connects PYL and CGL. That is, the sums of geometric series are not themselves geometric series, except for trivial cases.

To connect the two approaches it is useful to first state the CGL model and how it differs from the PYL model. Under CGL, the gain term is no longer $\frac{1}{t}$ as under decreasing-gain learning (DGL); instead, it is equal to some constant value, $\gamma \in (0, 1)$. Using our stochastic gradient learning approach, the general formulation of CGL for the linear model is:

$$a_{i,t} = a_{i,t-1} + \gamma \left[x_t - a_{i,t-1} - b' w_{t-1}\right] \tag{13}$$

One way to compare the two models is via a simulation exercise, to examine if there are any differences in the moments of the data generated by the PYL and CGL systems. As the CGL model has a free parameter for the gain, there is no natural correspondence between the PYL and the CGL as the PYL's gain sequence is pinned down by the birth and death rates. Hence, we need to make an assumption about the relationship between the birth and death rates and the CGL model in order to make any comparisons.

Here for comparison purposes, we assume that the weight on the first term is equal in both models, that is, $\gamma = S_1 = \frac{-g \ln\left(\frac{g}{1+g-q}\right)}{1-q}$. We believe this is the correct normalization to compare across the two models as this allows us to cleanly observe differences in the distribution of weights on past terms. In the Appendix, for robustness, we also report on a curve-fitting method as an alternative approach to comparing the two models. Using that approach, we also look at the differences in the weights between the two models.

Since most of the learning literature in macroeconomics uses data of a quarterly frequency, we specify time periods as quarters. Specifically, for the simulations we can assume $\gamma = \frac{-g \ln\left(\frac{g}{1+g-q}\right)}{1-q} = 0.018$ which is based on $g = 0.0031$ and $q = 0.0032$ which corresponds to the U.S. demographic data in 2019 at an annual birth rate of 1.24% and an average life expectancy of 72 years, which is 288 periods (quarters), which is a proxy for the death rate. The data are taken from the World Bank and the World Health Organization (WHO). In the Appendix, we provide a data description section where we specify exactly where the data comes from.

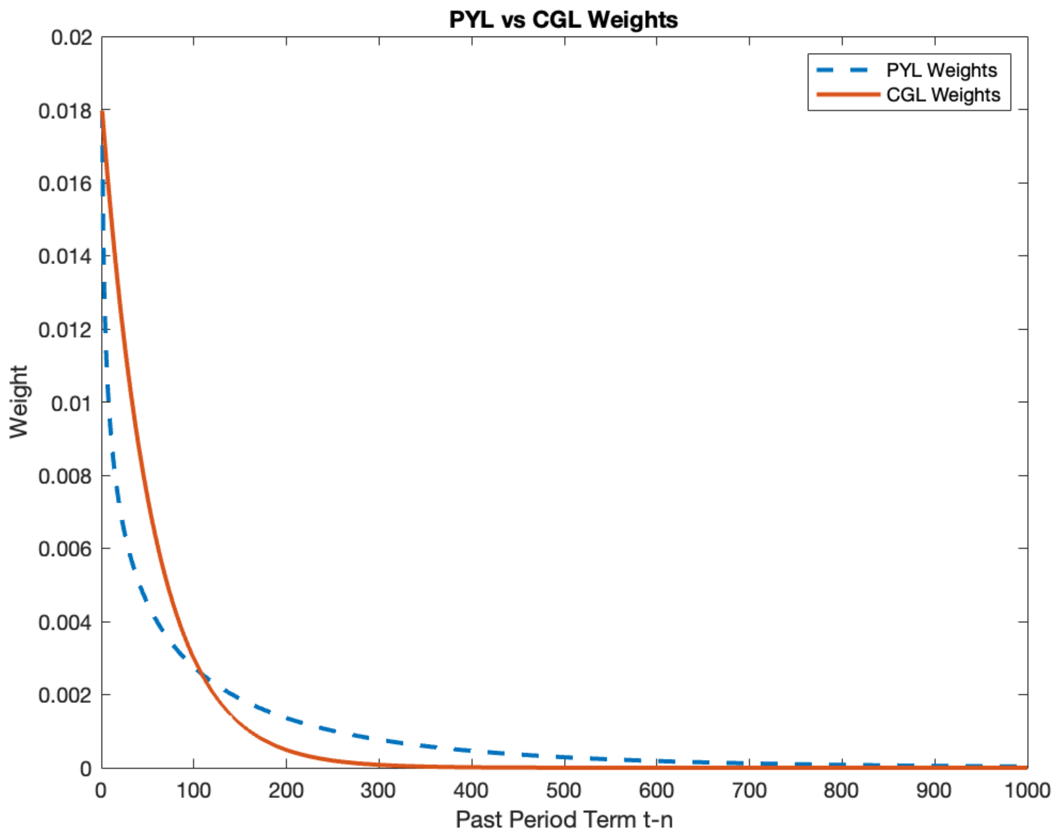

Figure 1: Differences in Weights Between PYL and CGL with $g = 0.0031$ and $q = 0.0032$

First, we show a graph that compares both models. Figure 1 reveals the weights that PYL and CGL place on past period residualized means $x_{t-n} - b'w_{t-n-1}$ given the g, q provided above. For PYL, the weights on past $x_{t-n} - b'w_{t-n-1}$ are as given by:[10]

$$x_t^e = b'w_{t-1} + S_1(g,q)(x_{t-1} - b'w_{t-2}) + S_2(g,q)(x_{t-2} - b'w_{t-3}) + ...$$

For CGL, the weights on past $x_{t-n} - b'w_{t-n-1}$ are given by:

$$x_t^e = b'w_{t-1} + \gamma(x_{t-1} - b'w_{t-2}) + \gamma(1-\gamma)(x_{t-2} - b'w_{t-3}) + ...$$

[10]Note here because $g \neq q$, we cannot characterize the $S_i(g,q)$ in terms of Lerch transcendents via the main results. When $g = q$, the graph is similar to Figure 1 which can be characterized exactly by the Lerch transcendent terms.

where here we need to normalize $\gamma = S_1$ as stated above. Note that Figure 1 is not a simulation but rather the implied weights of these two models given a set of parameters g, q.

As this figure reveals, the PYL model puts less weight on more recent observations than does the CGL model and PYL puts more weight on older observations than CGL. Because the curves in Figure 1 are distributions and integrate to 1, these differences between the two curves will be equal and offset one another. As the PYL model is an AR(∞) in expectations while the CGL model is an AR(1) in expectations, it is difficult to get closed-form expressions for the learning process for the PYL model. Therefore we adopt a computational strategy to see whether and how the two models differ.

4.1 Simulation

We again follow the normalization that the first weight on past $x_{t-1} - b' w_{t-2}$ under CGL will equal the first PYL weight. That is, given birth and death rates g, q, the equivalent CGL gain parameter would be:

$$\gamma_{CGL} = S_1 = \frac{-g \ln\left(\frac{g}{1+g-q}\right)}{1-q}$$

Notice that under this normalization, the CGL model is also pinned down by the birth and death rates. Another interpretation is that the CGL model is "approximately" the PYL model in the first lag term, which is the approach we will take below.

We simulate a simple linear asset pricing version of Equation (1) with AR(1) dividends under both CGL and PYL with different demographics (i.e., values for g). In the asset pricing interpretation, $x_t = p_t, x_t^e = p_{t+1}^e, b = R^{-1}, a = R^{-1}\mu, c = R^{-1}\rho$, and $W_{t-1} = D_t$.

$$p_t = R^{-1}(p_{t+1}^e + \mu + \rho D_t)$$

where p_t is the price at time t, p_{t+1}^e is the expected price at time $t+1$ formed at t, a is the constant term, ρ is the AR(1) term for dividends, and D_t is dividends at time t. Dividends follow:

$$D_t = \mu + \rho D_{t-1} + \epsilon_t$$

where $\mu \geq 0$ and $\epsilon_t \sim N(0, \sigma_D^2)$.

Expectations follow PYL learning where the PLM is:

$$p_t^e = a_{t-1} + bD_{t-1}$$
$$a_{t-1} = \sum_{i=1}^{\infty} q(1-q)^{i-1}\Phi(1-q,1,i)(p_{t-i} - b'D_{t-i-1})$$
$$b = \frac{\rho}{R-1}$$

To parameterize the model, we choose the following: $\mu = 1$, $R = 1.05$, $\rho = 0.9$, $\sigma_d^2 = 0.005$, which are used in Hommes and Zhu (2014).[11] We vary the birth rate, g, between 0.002 to 0.05, which corresponds to quarterly birth rates between 0.2% to 5%. As there is not much variation in life expectancy in OECD countries, we fix the death rate, q, in the simulations to 0.003, which corresponds closely to the US life expectancy of 72 years.[12] In order to calibrate the life expectancy, we can set q such that the average life expectancy in the model is $\frac{1}{q}$.

We run Monte Carlo simulations with 25,000 periods. We initialize each simulation with 5,000 initial observations (corresponding to the data that would be used by the oldest agent). We then perform 10,000 iterations of burn-ins, and 1,000 total runs per birth rate, g. We initialize the model at the implied REE of the models. Note that the reason we have 5,000 initial observations for the PYL model is that the model has to be initialized for the same number of periods as there are cohorts, such that the initial old cohorts are able to form beliefs as the model runs.

Table 1 provides a summary of our Monte Carlo simulation results using the CGL and PYL models for both prices, p_t, and expected prices p_{t+1}^e for different birth rates, g. Recall that under our normalization, the first weight in the CGL model is also pinned down by these same birth rates.

[11] The standard deviation of the error term is kept fairly low to ensure there are no escape paths in the simulations.

[12] Surprisingly, death rates are not as important as birth rates in the model, so our results are robust to variations in q. We will elaborate on this result in the next section.

Table 1: Simulation Results for the Asset Pricing Model by Birth Rates

		CGL		**PYL**	
g-value	Moment	Price	Expected Price	Price	Expected Price
g = 0.002	Mean	200	200	199.84	199.84
	SD	0.743	0.598	0.537	0.279
	Corr	0.398	0.342	0.291	0.167
g = 0.0035	Mean	200	200	200	200
	SD	0.916	0.804	0.626	0.418
	Corr	0.475	0.438	0.339	0.246
g = 0.005	Mean	199.99	199.99	200	200
	SD	1.049	0.955	0.716	0.539
	Corr	0.529	0.502	0.385	0.312
g = 0.01	Mean	199.99	199.99	199.99	199.99
	SD	1.346	1.276	0.932	0.801
	Corr	0.626	0.612	0.482	0.437
g = 0.02	Mean	199.99	199.99	199.99	199.99
	SD	1.733	1.682	1.258	1.163
	Corr	0.721	0.715	0.601	0.577
g = 0.05	Mean	199.99	199.99	199.98	199.98
	SD	2.347	2.314	1.856	1.795
	Corr	0.817	0.816	0.745	0.738

Notes: SD is standard deviation, Corr refers to the first order autocorrelation of the column variable

As Table 1 reveals, PYL does not change mean prices or expected prices relative to CGL; instead, PYL primarily affects the standard deviation, which is lower under PYL. While this difference is not very large, the reason for this difference is as follows. Under PYL, a part of the history of observations that agents use to update their expectations is removed, in particular the history that occurred before they were born. That is, under CGL, each agent while adding more weight to current observations is still weighting each past term. In contrast, under PYL, for each agent, some of the past observations are not weighted at all. When aggregated, the PYL weights will be missing parts of the history and this rather mechanically leads to lower volatility. As the missing history is mainly in the tail terms, this does not contribute much quantitatively to the overall volatility.

To understand this result more formally, let us think about the role of the constant-gain parameter γ and the variance or standard deviation of x_t. In a simple linear model with no regressors which is Equation (1), where $\delta = 0$, the relationship between γ and the variance of x_t is $V(x_t) = \frac{1+(1-\gamma)(1-2\beta)}{1+(1-\gamma)(1-2\beta)-\gamma\beta^2}\sigma_\epsilon^2$. Here, $\frac{\partial V(x_t)}{\partial \gamma} > 0$ so a higher gain implies a higher variance for x_t. Note that this relationship is also true more generally for the case with exogenous regressors but here we can see the relationship more clearly in closed form.

Then, we can do a variance decomposition that is implied by our simulations above as follows. For a given g, q there is PYL model that generates a variance for prices. This model has a corresponding CGL model whose first weight is pinned down by the same p, q used in the PYL model. Then the variance of the CGL model equals the variance of the PYL model plus the variance of the missing history component. Because of this missing history component, the PYL model has a lower relative CGL gain, and therefore a slightly lower variance as shown in Table 1.

Overall, the general pattern in the simulations is that both price and expected price variation are slightly lower under PYL versus CGL with accompanying decreases in the first-order autocorrelation of prices and expected prices. One interpretation is that because the DGP relies only on the one-step ahead expectations, the model reverts to the steady state quickly so any under or over-reactions are quickly corrected. In fact it turns out that PYL has a lower autocorrelation for higher-order terms as well. We suspect this is related to the fact that PYL has a lower relative variance versus the CGL and the linearity of the model with respect to expectations.

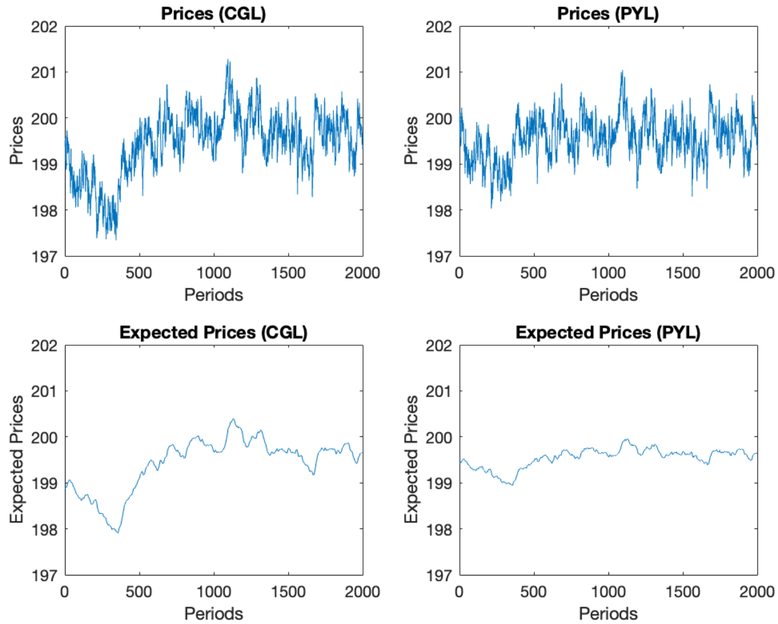

Figure 2: Simulations with CGL and PYL with $g = 0.0035$

We also show Figure 2 from two individual simulations with $g = 0.0035$ to visually compare the differences using a representative sub-sample of a simulation run. Note that this is a particular simulation of a smaller length for expositional purposes. As we can see again, the models look quite similar with slightly lower volatility in expected prices under the mechanism explained above. Here, the volatility of prices in the CGL model is 0.864 and the volatility of prices in the PYL model is 0.582. The volatility of expected prices in the CGL model is 0.756 and the volatility of expected prices in the PYL model is 0.371. As this is a particular simulation, the numbers are slightly different than the averaged values from the Monte Carlo simulations in Table 1.

5 Approximation Analysis

In this section, we provide an analytical approximation for the CGL model using the PYL model. In the Appendix, for robustness, we provide another approximation approach using a curve-fitting algorithm to find the best-fitting PYL curve that matches the CGL curve conditional on a fixed birth rate g.

Recall again from Equation (10), that there is no exact correspondence between PYL and CGL as sums of geometric series do not lead to geometric series. We show above that the numerical properties of the models are similar in the moments. As in the numerical simulations shown in Section 4, we can find an approximate CGL formulation, given our PYL model as follows:

$$a_t = a_{t-1} + \frac{-g \ln\left(\frac{g}{1+g-q}\right)}{1-q} \left[x_t - a_{t-1} - b' w_{t-1}\right]$$

Then, the equivalent CGL gain parameter γ_{CGL} is pinned down by the demographics of the model g, q. Hence, depending on the questions that the modeler is interested in such as questions about first and second moments, an approximate CGL model can be microfounded and the gain parameter can be pinned down via the demographic data.

5.1 Implied Gains Over the Cross-Section

Here we show a cross-section of birth and death rates around the world to see what the implied constant gain is for a given g, q pair. We choose the U.S., Japan, China, India, and Nigeria to cover a wide range of birth and death rates.

Table 2: Cross-Sectional Variation in Birth and Death Rates

Country	Birth Rate g	Death Rate q	Implied Gain
US	0.31%	0.320%	0.0180
Japan	0.20%	0.290%	0.0125
China	0.30%	0.323%	0.0175
India	0.47%	0.353%	0.0253
Nigeria	0.88%	0.399%	0.0418

The values in Table 2 are quarterly rates. Birth rates are from the WHO and death rates are imputed from life expectancy from the World Bank.[13] Here we find fairly large variations in implied gains from 0.0125 - 0.0418, with the main result that lower birth rates lead to a lower population with less young which leads to a lower gain and a higher birth rate leads to more young, and hence a higher gain. We can see that there is not much variation in life expectancy in developed countries and that all the variation in implied gains is due to variations in the birth rate.

[13] Our results are also robust to a variety of death rates such as mortality rates from the WHO, as death rates do not have a significant quantitative impact on the results.

Interestingly, our model results for developed countries are closely aligned with gain values that are found in empirical studies (0.01 - 0.04) (Orphanides and Williams (2005), Branch and Evans (2006), Milani (2007), Slobodyan and Wouters (2012), Malmendier and Nagel (2016), Berardi and Galimberti (2017), and Eusepi and Preston (2018)). Strikingly, calibrating the model to U.S. demographics implies a gain of 0.018 which is almost identical to that found by Milani (2007) of 0.0183 and identical to Malmendier and Nagel (2016) of 0.018.

Surprisingly, death rates do not matter that much in our model. If we look at the CGL formulation: $\gamma_{CGL} = \frac{-g\ln\left(\frac{g}{1+g-q}\right)}{1-q}$, we find that the gain is almost entirely determined by g. Moreover, for empirically relevant q, which are small and close to g, this implies that $1+g-q \approx 1$ and $1-q \approx 1$ and hence, $\frac{\partial \gamma_{CGL}}{\partial q} \approx 0$. On further investigation, we find that because of the structure of the perpetual youth model, the distribution is skewed towards the young. Thus, changes in the death rate q do not make a large impact on the distribution and hence the weight placed on the gain. As such, the model implies that variations in the birth rate will mainly determine the variations in the gain.

5.2 Implied Gains Over Time

Our model also makes predictions for implied gains due to demographic impacts over time, so we can also look at how these gains have evolved over time. We focus on the cases of the U.S. and Japan. Table 3 lists decade-average birth rates for these two countries. Since death rates do not matter much in the model, we use the 2019 life expectancy from the World Bank throughout in the tables below.

Table 3: Birth Rates Over Time and Implied Gains for the U.S. and Japan

	U.S.		Japan	
Decade	Birth Rate g	Implied Gain	Birth Rate g	Implied Gain
1950s	0.599%	0.0308	0.552%	0.0288
1960s	0.499%	0.0265	0.444%	0.0241
1970s	0.388%	0.0216	0.431%	0.0236
1980s	0.385%	0.0215	0.310%	0.018
1990s	0.375%	0.0210	0.245%	0.0148
2000s	0.349%	0.0198	0.222%	0.0136
2010s	0.312%	0.0181	0.202%	0.0126

The general pattern in developed countries such as the U.S. and Japan is slowly declining birth rates. This implies that the model-implied gains are decreasing over time. For the U.S., birth rates peaked at a quarterly rate of 0.599% around the baby boom period with an implied gain of 0.0308, which decreased over time to 0.0181 in the 2010s. Japan, on the other hand, had a much sharper decline in birth rates in the 80s and 90s, which eventually led to implied gains of 0.0126 in the 2010s.

Here, our model may help to both consolidate and explain some of the findings in the literature. In particular, the time-variation in gains that we see in Branch and Evans (2006) and Slobodyan and Wouters (2012) when looking at different samples and time periods in the data might be partially explained by declining birth rates over time. Cross-country differences in constant-gain estimates for inflationary expectations among South American countries as reported by Sargent et al. (2009) could also partly reflect differences in birth and death rates across these countries.

Berardi and Galimberti (2017) make an important point that the specification of the model matters for the constant-gain estimates. Here we try to formalize one set of structural parameters that may help explain why model specifications may lead to variations in the implied gains. Again, we do not claim that our model mechanism is the only one happening in the data, but the reality is most likely a mix of many mechanisms found in the literature, with an important one possibly being CGL algorithms as optimal responses to parameter drift (Evans et al. (2010)).

Ultimately, our model provides both a method of disciplining a free parameter for CGL and an interpretation of both the time-varying gains and variations in the cross-section of gains we see in the data. The story that the model tells is one where declining birth rates have decreased gains over time in developed countries and one where countries with low birth rates (e.g., Japan) are expected to have lower gains than countries with higher birth rates (e.g., India). Our model provides a potential explanation for the difficulty of moving anchored inflation expectations (Coibion and Gorodnichenko (2015)) in countries like Japan, which also has among the lowest birth rates relative to other OECD countries.

6 Alternative Learning Rules

Thus far we have considered the case where agents learn only from data over their own lifetimes and using the standard DGL approach where the individual gain is the OLS gain $\frac{1}{i}$. It is of interest to consider cases where individual agents may use other rules that consider data before they were born or that depart from the decreasing-gain specification that we use. Therefore, in this section, we consider three alternative approaches to modeling learning in our PYL model and we relate the effective gains from these alternative learning models to CGL learning. For tractability, we will focus on the case again where the birth rate equals the death rate $g = q$.

6.1 Additional Data

The first modification we consider is to allow agents to use data not just within their own life horizon. In the PYL setting, we can think of this case as "learning from parents". Formally, we allow agents to include h additional periods into their information sets, by putting weight $\frac{1}{i+h}$ onto all the past observations, which is equivalent to the formulation being:

$$\gamma^h_{CGL} = S_{1+h} = q(1-q)^h \Phi(1-q, 1, 1+h)$$

Another way to think about this is that adding h data points to an agent's information set is the same as having more agents of age $i+h$ in the PYL model.[14] Then, once we specify h, we can calculate the implied effective constant gain corresponding to the Lerch transcendent term S_{1+h} further in the sequence. In general, the corresponding weight will be lower than for the S_1 case.

As illustrative examples, we consider the cases where $h = 0, 2, 4, 6, 8,$ and 10, where $h = 0$, is just our benchmark PYL model.[15] Using the parameterization above where $g = q = 0.0031$, we get the following implied effective constant gains:

[14]More specifically since every cohort includes h more data points into their information set, this learning rule is like "shifting" the Lerch transcendent terms, S_i by h terms.

[15]In the Online Appendix, we provide a larger parameterization table for all three learning rules along with varying them by the demographic parameter q.

Table 4: Alternative Learning Rule: Additional Data

Additional Data (h)	Effective CGL Gain (γ^h_{CGL})
h = 0 (PYL)	0.018
h = 2	0.013
h = 4	0.011
h = 6	0.010
h = 8	0.009
h = 10	0.009

As Table 4 reveals, as h increases, the effective gain decreases. Intuitively,

the young effectively weigh "parental" data the same as if they had lived through it themselves. As an example, if $h = 4$, then $S_{1+h} = S_5$ would be the weight via the age 5 cohort. Hence, this learning rule generally implies a lower effective constant gain than our benchmark PYL model.

6.2 General Decreasing-gain Learning (DGL)

We next consider the case of a more general decreasing-gain function, that is, agents use the following gain $\frac{1}{i^\alpha}$ for $\alpha > 0$. Note that $\alpha = 1$ corresponds to the least-squares gain case we use in our PYL model $\frac{1}{i}$.

We show that the learning rule under $\frac{1}{i^\alpha}$ leads to another transcendental form known as a polylogarithm function, which reduces to the Lerch transcendent when $\alpha = 1$.[16] To determine the implied effective constant gain, we can just utilize the same approach and replace the least-squares gain $\frac{1}{i}$ with the general decreasing gain $\frac{1}{i^\alpha}$.

Then this implies, in the case of $g = q$:

$$S_1^\alpha = \sum_{i=1}^{\infty} \frac{q(1-q)^{i-1}}{i^\alpha} = \frac{qLi_\alpha(1-q)}{1-q}$$

where:

$$Li_\alpha(1-q) = \sum_{k=1}^{\infty} \frac{(1-q)^k}{k^\alpha} = (1-q) + \frac{(1-q)^2}{2^\alpha} + \frac{(1-q)^3}{3^\alpha} + ...$$

Here $Li_\alpha(\cdot)$ is the polylogarithm function. One downside of the polylogarithm function is

[16] More specifically, it reduces to our formulation of the Lerch transcendent, namely, $\Phi(1-q, 1, i)$.

that, unlike the Lerch Transcendent, the polylogarithm function does not have a simple analytical form that allows us to determine general S_i^α terms.

We report numerical values for the implied effective constant gain and compare these with our results under the same constant gain approximation as we do in our PYL model as follows:

$$\gamma_{CGL}^\alpha = S_1^\alpha = \sum_{i=1}^{\infty} \frac{q(1-q)^{i-1}}{i^\alpha} = \frac{qLi_\alpha(1-q)}{1-q}$$

Here we need to set α to specific values. We choose the following: $\alpha = 0.5, 0.75, 1, 1.25, 1.5, 1.75, 2$. Note that $\alpha = 1$ is just the benchmark PYL model. Then we find that:

Table 5: Alternative Learning Rule: General DGL

General DGL (α)	Effective CGL Gain (γ_{CGL}^α)
$\alpha = 0.5$	0.092
$\alpha = 0.75$	0.038
$\alpha = 1$ (PYL)	0.018
$\alpha = 1.25$	0.011
$\alpha = 1.5$	0.008
$\alpha = 1.75$	0.006
$\alpha = 2$	0.005

Table 5 shows the different implied effective constant gains for the general DGL model. Here we find that for α values under 1, the effective gains are larger than those found in our main results where $\gamma_{PYL} = S_1^1 = 0.018$. This is because as $\alpha \to 0$, the generalized DGL model approaches naïve learning as agents place higher and higher weight on the most recent observation.

For α values greater than 1, the implied effective constant gains are decreasing. In particular, as α increases, agents are putting less and less weight on new observations relative to the least-squares gain case as they are placing weight $\frac{1}{i^\alpha}$ on the value. For α values in the neighborhood of 1, our results are robust to the use of the general decreasing-gain learning rule, with α values between 0.75 - 1.25 yielding effective constant gains within the range of empirical estimates.

6.3 Biased Weights (Malmendier and Nagel, 2016)

Finally, we consider one version of the learning model used in Malmendier and Nagel (2016) (MN) and Malmendier et al. (2020), where the weight that agents place on data is given by:

$$w(k, \lambda, age) = \frac{(age + 1 + k)^{\lambda}}{\sum_{k'=0}^{age}(age + 1 - k')^{\lambda}}$$

Here $w(\cdot)$ is the weight on past observations, k is the period, k' is the index, λ is the degree of recency bias, and *age* is the age of the agent. In the MN learning rule, $\lambda = 1$, while in our benchmark PYL model, $\lambda = 0$, which is the DGL case.[17]

As shown below, recency bias results in a learning rule where agents linearly weight the most recent observations much higher than in our approach. As an example, cohorts of age 3 use the following scheme:

$$\begin{aligned} a^{e}_{i,t} &= w(0,1,2)(x_t - b'w_{t-1}) + w(1,1,2)(x_{t-1} - b'w_{t-2}) + w(2,1,2)(x_{t-2} - b'w_{t-3}) \\ &= \left(\frac{3^{\lambda}}{1 + 2^{\lambda} + 3^{\lambda}}\right)(x_t - b'w_{t-1}) + \left(\frac{2^{\lambda}}{1 + 2^{\lambda} + 3^{\lambda}}\right)(x_{t-1} - b'w_{t-2}) + \\ &\left(\frac{1}{1 + 2^{\lambda} + 3^{\lambda}}\right)(x_{t-2} - b'w_{t-3}) \end{aligned}$$

where with $\lambda = 1$ (as in MN), we call this particular learning rule "biased weights" as agents exhibit *recency bias* even within their personal histories. Note that in general, with $\lambda = 1$ the denominator will be determined by the formula for the sum of the first n natural numbers:

$$a_i = \frac{1}{2} i * (i + 1)$$

Using the fact that the most recent observation receives weight i and every subsequent observation $i - j$ receives weight $i - j$, we can characterize the weight S_i on each past observation $x_t - b'w_{t-1}$ similar to our approach in the benchmark PYL model. More importantly, as we will use the first term as our approximation, the first observation will receive weight i over the sum of the first i natural numbers.

The formula for the implied effective constant gain under biased weights (MN) is then given by Proposition 2:

Proposition 2. *Under the biased weights learning rule (MN), the weight on the first term*

[17] Note that this leads to the least-squares gain $\frac{1}{i}$ when $\lambda = 0$ because each individual term is then reduced to 1.

S_1^{MN} *is characterized by the infinite sum of the following series, which is:*

$$\gamma_{CGL}^{MN} = S_1^{MN} = \sum_{i=1}^{\infty} q(1-q)^{i-1} * \frac{i}{\frac{1}{2}i(i+1)} = \frac{2q(q - \ln q - 1)}{(1-q)^2}$$

We prove Proposition 2 in the Appendix.

We find that the implied effective constant gain under the biased weights learning rule is: $\gamma_{CGL}^{MN} = 0.031$, which is higher than what we found using the benchmark PYL model with $\gamma_{PYL} = 0.018$. Still, it is within the range of empirical estimates of the constant gain found in the literature (recall this range among key papers was found to be 0.01 - 0.04). Moreover, the fact that the biased weights imply a higher gain is intuitive given that the MN learning rule places much more weight on the most recent observation than our benchmark PYL model which results in a relatively higher implied effective constant gain.

Nevertheless, despite these different ways of modeling individual learning, we continue to find that our PYL approach delivers implied effective constant gains that are empirically plausible. More importantly, the method that we introduce with our benchmark PYL model is one that can be applied more broadly to a variety of learning rules in order to connect it to a constant-gain learning rule. One can imagine that in reality a variety of these different learning approaches are operating simultaneously and thus the *average* of these approaches would yield something close to the implied effective constant gain via our benchmark PYL model.

7 Conclusion

We provide an analytical characterization of the perpetual youth model with learning-from-experience (LfE) in a linear expectations model. For our baseline case of decreasing-gain learning, we find that the PYL differs from CGL in that the formulation is described by the Lerch transcendent rather than geometrically declining weights in past data. One benefit of the characterization is that it ties expectations to the birth and death rates while CGL depends on a free parameter for the gain. PYL also nests two well-known learning models: DGL and naïve learning as special cases.

We find that with simulations under a standard asset pricing model with AR(1) dividends, the moments under CGL and PYL are approximately the same, that is, PYL has slightly lower standard deviations under PYL but the means are near identical. We do Monte Carlo simulations for a range of values and find that the pattern persists. We conclude that PYL can approximate a CGL model in the first and second moments while also microfounding the gain parameter to the demographics of the model. We find that when calibrating the gain to realistic demographics, the implied gain is reasonably within values found in empirical studies and identical to those found in some key papers. We further show that our PYL approach is robust to the use of some alternative learning models in the sense that the effective gain is generally within the range of values that have been estimated using constant-gain learning approaches.

One may wonder why the individual agents in our model use various decreasing-gain algorithms despite observing that the world in which they operate is non-stationary. The choice between constant and decreasing-gain algorithms depends on the specific characteristics of the economic environment being modeled. In rapidly changing environments, constant-gain algorithms may be more appropriate due to their flexibility. However, in "normal times" where changes are less frequent or more gradual, decreasing-gain algorithms are a reasonable choice for their stability and reduced sensitivity to recent shocks or noise. Still, future research using our approach might consider the case where the individuals themselves use constant-gain learning, or perhaps some endogenous switching between decreasing and constant-gain learning as in Cho et al. (2002) and Marcet and Nicolini (2003).

Other extensions to our approach might explore in further detail the microfoundations for the appropriate number of periods that agents use within their information set. One could also extend our approach to least squares learning, and learning about coefficients on the regressors as well. We leave these extensions to future research.

References

Abramowitz, Milton and Irene A Stegun (1964) *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, 55: National Bureau of Standards, Dover.

Adam, Klaus and Albert Marcet (2011) "Internal rationality, imperfect market knowledge and asset prices," *Journal of Economic Theory*, 146 (3), 1224–1252.

Adam, Klaus, Albert Marcet, and Johannes Beutel (2017) "Stock price booms and expected capital gains," *American Economic Review*, 107 (8), 2352–2408.

Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini (2016) "Stock market volatility and learning," *Journal of Finance*, 71 (1), 33–82.

Berardi, Michele (2020) "A probabilistic interpretation of the constant gain learning algorithm," *Bulletin of Economic Research*, 72 (4), 393–403.

Berardi, Michele and Jaqueson K Galimberti (2017) "Empirical calibration of adaptive learning," *Journal of Economic Behavior & Organization*, 144, 219–237.

Branch, William A and George W Evans (2006) "A simple recursive forecasting model," *Economics Letters*, 91 (2), 158–166.

Bray, Margaret M and Nathan E Savin (1986) "Rational expectations equilibria, learning, and model specification," *Econometrica*, 1129–1160.

Cho, In-Koo, Noah Williams, and Thomas J Sargent (2002) "Escaping nash inflation," *Review of Economic Studies*, 69 (1), 1–40.

Coibion, Olivier and Yuriy Gorodnichenko (2015) "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 105 (8), 2644–2678.

Eusepi, Stefano and Bruce Preston (2011) "Expectations, learning, and business cycle fluctuations," *American Economic Review*, 101 (6), 2844–2872.

——— (2018) "The science of monetary policy: An imperfect knowledge perspective," *Journal of Economic Literature*, 56 (1), 3–59.

Evans, George W and Seppo Honkapohja (2001) *Learning and expectations in macroeconomics*: Princeton University Press.

——— (2009) "Learning and macroeconomics," *Annual Review of Econonmics*, 1 (1), 421–449.

Evans, George W, Seppo Honkapohja, and Noah Williams (2010) "Generalized stochastic gradient learning," *International Economic Review*, 51 (1), 237–262.

Hertwig, Ralph, Greg Barron, Elke U Weber, and Ido Erev (2004) "Decisions from experience and the effect of rare events in risky choice," *Psychological Science*, 15 (8), 534–539.

Hommes, Cars and Mei Zhu (2014) "Behavioral learning equilibria," *Journal of Economic Theory*, 150, 778–814.

Kreps, David M (1998) "Anticipated utility and dynamic choice," *Econometric Society Monographs*, 29, 242–274.

Malmendier, Ulrike (2021) "Experience effects in finance: Foundations, applications, and future directions," *Review of Finance*, 25 (5), 1339–1363.

Malmendier, Ulrike and Stefan Nagel (2011) "Depression babies: Do macroeconomic experiences affect risk taking?" *Quarterly Journal of Economics*, 126 (1), 373–416.

——— (2016) "Learning from inflation experiences," *Quarterly Journal of Economics*, 131 (1), 53–87.

Malmendier, Ulrike, Demian Pouzo, and Victoria Vanasco (2020) "Investor experiences and financial market dynamics," *Journal of Financial Economics*, 136 (3), 597–622.

Marcet, Albert and Juan P Nicolini (2003) "Recurrent hyperinflations and learning," *American Economic Review*, 93 (5), 1476–1498.

Marcet, Albert and Thomas J Sargent (1989) "Convergence of least squares learning mechanisms in self-referential linear stochastic models," *Journal of Economic Theory*, 48 (2), 337–368.

Milani, Fabio (2007) "Expectations, learning and macroeconomic persistence," *Journal of Monetary Economics*, 54 (7), 2065–2082.

——— (2008) "Learning, monetary policy rules, and macroeconomic stability," *Journal of Economic Dynamics and Control*, 32 (10), 3148–3165.

——— (2014) "Learning and time-varying macroeconomic volatility," *Journal of Economic Dynamics and Control*, 47, 94–114.

Nakov, Anton and Galo Nuño (2015) "Learning from experience in the stock market," *Journal of Economic Dynamics and Control*, 52, 224–239.

Orphanides, Athanasios and John C Williams (2005) "The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations," *Journal of Economic dynamics and control*, 29 (11), 1927–1950.

Sargent, Thomas J and Noah Williams (2005) "Impacts of priors on convergence and escapes from Nash inflation," *Review of Economic Dynamics*, 8 (2), 360–391.

Sargent, Thomas, Noah Williams, and Tao Zha (2009) "The conquest of South American inflation," *Journal of Political Economy*, 117 (2), 211–256.

Shin, Michael (2021) "Subjective expectations, experiences, and stock market participation: Evidence from the lab," *Journal of Economic Behavior & Organization*, 186, 672–689.

Slobodyan, Sergey and Raf Wouters (2012) "Learning in an estimated medium-scale DSGE model," *Journal of Economic Dynamics and control*, 36 (1), 26–46.

Weber, Elke U, Ulf Böckenholt, Denis J Hilton, and Brian Wallace (1993) "Determinants of diagnostic hypothesis generation: effects of information, base rates, and experience.," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19 (5), 1151–1164.

Williams, Noah (2019) "Escape dynamics in learning models," *Review of Economic Studies*, 86 (2), 882–912.

A Appendix

Data Description

Here we specify the exact sources for the demographic data used for birth and death rates. The most common resources are from the World Health Organization (WHO), CIA Factbook, and the World Bank. Birth and death rate measurements can also vary depending on the year and gender.

For birth rates, we use the CIA World Factbook based on the 2023 population estimates and compare the average annual number of births during a year per 1000 births at midyear. This is sometimes known as a crude birth rate. As this rate is annual, we take the quarterly rate using an effective rate calculation to find the effective quarterly rate that implies the effective annual rate as follows:

$$g_{annual} = (1 + g_{quarter})^4 - 1$$

Nevertheless, because birth rates are small, rounding to 4 significant digits will make the rough calculation of dividing the annual rate by 4 essentially equivalent.

For death rates, there are two main alternatives. We can either take implied mortality rates which calculate the number of deaths over a period divided by the person-years lived by the population over that period, sometimes called the crude death rate, or the life expectancy which states how long a given person can be expected to live up to. As death rates are not important quantitatively for our model, our calibrations and simulations are robust to either specification. Here we choose life expectancy as our proxy for the death rate as crude death rates can be misleading as developing countries have far more young than old versus developed countries, and also may have much higher age-specific mortality rates while having lower crude mortality rates. Hence, for life expectancy, we use data from the Global Health Observatory from the World Health Organization (WHO) on overall life expectancy for both genders in 2019.

For historical birth rates in the U.S. and Japan, we use collated data from the data aggregation site Macrotrends which sources the historical data from the United Nations World Population Prospects.

Optimal Fitted CGL

For robustness, we also do an alternative formulation for finding the CGL gain parameter γ that best fits the PYL formulation. Specifically, we set up the following minimization problem:

$$\arg\min_{\gamma} \sum_{i=1}^{\infty} \left[\gamma(1-\gamma)^{i-1} - S_i(g)\right]^2 \tag{14}$$

where, $S_i(g)$ is the weights on the past $x_{t-i} - b'w_{t-i-1}$ terms. That is, we minimize the squared distance between the two curves: CGL and PYL and find the γ that is optimal for this objective function.

In order to operationalize the minimization problem, we utilize the stationary, no growth population PYL and set $g = q$ which implies:

$$x_t^e = \sum_{i=1}^{\infty} g(1-g)^{i-1}\Phi(1-g, 1, i)(x_{t-i} - b'w_{t-i-1}) + b'w_{t-1}$$

where here since $g = q$, we use g instead of q. This is required as the Lerch transcendent has an analytical form that can be much more easily implemented computationally.[18] Since we are unable to optimize over the infinite sum computationally, to operationalize this, we set the upper limit of summation to 1000.[19] We do this for a grid of values for $g \in [0, 1]$.

[18]Since g has a much more significant impact on S_i than q, the stationary, no growth population PYL equation with g replacing q well approximates the growing population case.

[19]Because the tails of the series are all close to 0, the optimization results are not sensitive to larger choices for the upper limit of summation.

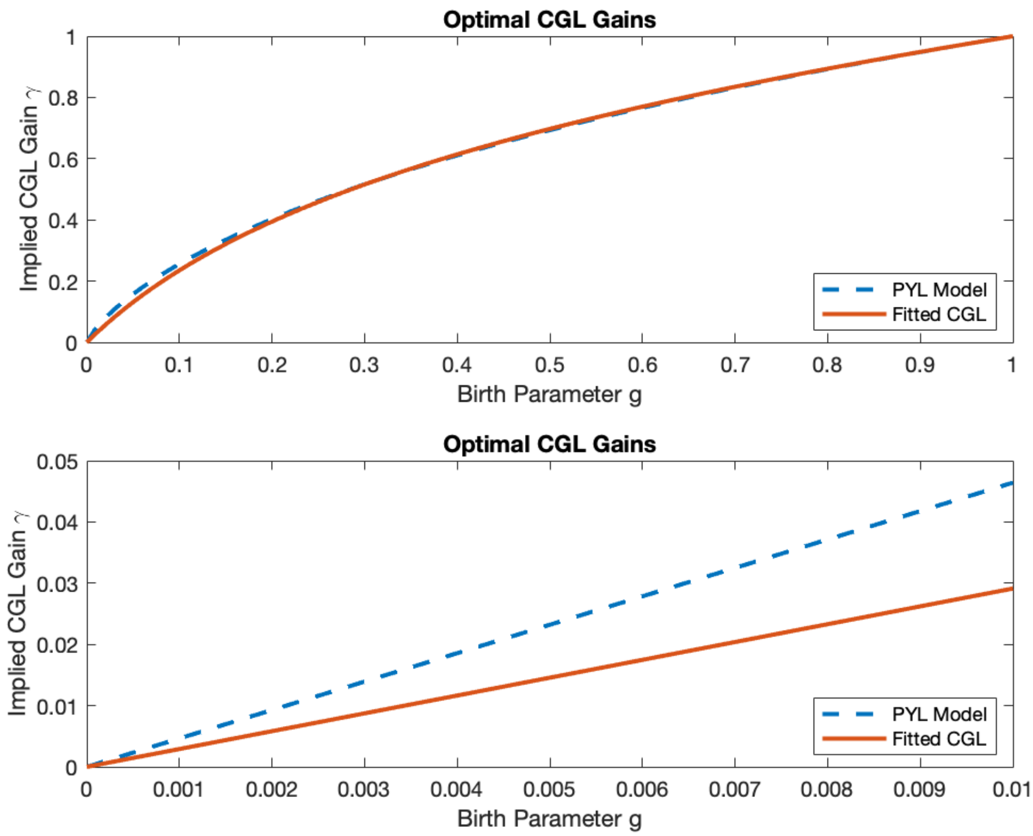

Figure 3: Implied gains γ and S_1 by Birth Rate, g, top panel $g = [0, 1]$ bottom panel $g = [0, 0.01]$

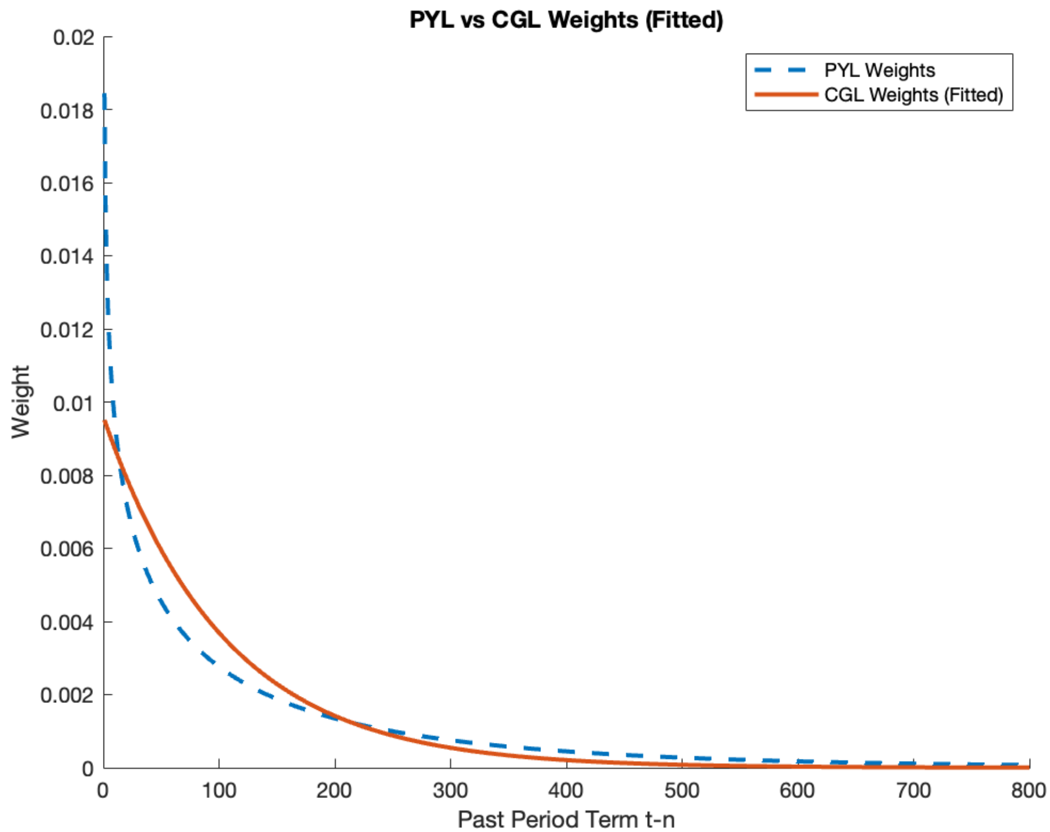

Figure 4: Fitted CGL Gain for g = 0.0031

In Figure 3, the top panel, we graph the optimal CGL gains γ obtained from Equation (14) for each value of g within the grid $[0,1]$. In Figure 3, the bottom panel, we show the results for a restricted domain $g \in [0, 0.01]$ which is more relevant for realistic birth rates. We also show the CGL gains γ found by setting $\gamma = S_1 = \frac{-g\ln\left(\frac{g}{1+g-q}\right)}{1-q}$ per our earlier approximation (PYL model) compared with those found with our optimization results in Equation (14) (Fitted CGL) in Figure 4. To better make sense of this result, we show the optimization results comparing the optimal fitted CGL versus its corresponding PYL curve for $g = 0.0031$ in Figure 4.

First, we find that for the demographics in the U.S., the implied gain parameter in our approximation of $\gamma = S_1$ overstates the implied gain with a value of $\gamma = 0.018$, with the curve fitting approach from Equation (14), which is around $\gamma = 0.01$. Next, we see that for most of the domain between $g \in [0,1]$, $S_1 = \gamma$ approximates the optimal curve fitted γ quite well. It seems that for the domain, we care about for empirical values between

$g \in [0, 0.01]$, both curves diverge with the $S_1 = \gamma$ approximation overstating the implied gain parameter found in the curve fitting exercise and start to converge again further along the curve. Nevertheless, our approach is robust to the curve-fitting approach with implied CGL gain parameters approximately 0.005 to 0.01 points lower than those done by our $\gamma = S_1$ approach depending on the range for g.

Proof of Proposition 1

Here we use well-known results about generating functions to prove the proposition. We first solve for $S_1 = \sum_{i=1}^{\infty} \frac{q(1-q)^{i-1}}{i}$. It is straightforward to confirm that S_1 converges for $g, q \in (0, 1)$. Note that S_1 looks of the form: $\sum_{i=1}^{\infty} \frac{x^i}{i}$ which is a geometric-harmonic series and is known to converge to $-\ln(1-x)$. Here $A(x) = \sum_{i=1}^{\infty} \frac{x^i}{i} = \sum_{i=1}^{\infty} \frac{x^{i-1} * x}{i}$. Taking the x out and dividing $A(x)$ implies $\sum_{i=1}^{\infty} \frac{x^{i-1}}{i} = \frac{A(x)}{x} = \frac{-\ln(1-x)}{x}$. Plug in $x = 1 - q$ implies $S_1 = \frac{-q \ln q}{1-q}$. For $S_1' = \sum_{i=1}^{\infty} \frac{g(1-q)^{i-1}}{(1+g-q)^i} * \frac{1}{i}$, let $x = \frac{1-q}{1+g-q}$ and apply the same steps as above. Then $S_1' = \frac{-g \ln(\frac{g}{1+g-q})}{1-q}$.

Proof of Proposition 2

Here we prove Proposition 2 which is the implied constant-gain formulation for the alternative learning rule for biased weights as in Malmendier and Nagel (2016). As before we use well-known results about generating functions. First we find that we can rearrange the sum as $\sum_{i=1}^{\infty} q(1-q)^{i-1} \frac{i}{\frac{1}{2}i(i+1)} = 2q \sum_{i=1}^{\infty} (1-q)^{i-1} \frac{1}{i+1}$. Then once we find the sum for $\sum_{i=1}^{\infty} \frac{x^i}{i+1}$, we can shift the index and multiply through with constants. Note that it is known that the following infinite sum has the following solution: $\sum_{i=1}^{\infty} \frac{x^i}{i+1} = \frac{-x-\ln(1-x)}{x}$. Then define $A(x) = \sum_{i=1}^{\infty} \frac{x^i}{i+1} = \sum_{i=1}^{\infty} \frac{x^{i-1} * x}{i+1}$. Taking the x out and dividing $A(x)$ implies that $\sum_{i=1}^{\infty} \frac{x^{i-1}}{i+1} = \frac{A(x)}{x} = \frac{-x-\ln(1-x)}{x} * \frac{1}{x} = \frac{-x-\ln(1-x)}{x^2}$. Then $2 * \frac{A(x)}{x} = 2 * \frac{-x-\ln(1-x)}{x} * \frac{1}{x} = \frac{-2x-2\ln(1-x)}{x^2}$. Then plugging in $x = 1 - q$, multiplying by q, and rearranging implies that this equals $\frac{2q(q-\ln q-1)}{(1-q)^2}$ and then this implies that $S_1^{MN} = \frac{2q(q-\ln q-1)}{(1-q)^2}$.

PYL with Learning about Exogenous Variables

Here we provide the full PYL formulation when learning about the b_{t-1} coefficients on the exogenous variables w_{t-1} for the case with $g = q$. Under this formulation, agents must now make estimates for both the constant term a_t and the coefficient b_t. Then x_t^e follows:

$$x_t^e = a_{t-1} + b_{t-1}w_{t-1}$$
$$a_{t-1} = \sum_{i=1}^{\infty} q(1-q)^{i-1}\Phi(1-q,1,i)(x_{t-i} - b'_{t-i}w_{t-i-1})$$
$$b_{t-1} = \sum_{i=1}^{\infty} \left\{\Pi_{i=1}^{\infty}(1 - qw'_{t-i-1}w_{t-i-1})\Phi(1-q,1,i)(x_{t-i} - a_{t-i-1})\right\}$$

Note that there are aspects of this problem that make it intractable. First, because a_{t-1} is a function of b_{t-1} and vice-versa, it is difficult to analyze the pair of functions. Second, the function for b_{t-1} is now an infinite product rather than a sum. Hence, the approach that is used in the body of the paper is more tractable for our purposes. Future work may want to explore this approach more in-depth.