Chile's Fiscal Rule as Social Insurance

Eduardo Engel
Universidad de Chile

Christopher Neilson
Yale University

Rodrigo Valdés
International Monetary Fund

Well before the Great Recession of 2009 put fiscal policy debates in the front burner, commodity-exporting countries had to deal with important fiscal policy dilemmas stemming from revenue volatility and eventual depletion. Chilean policymakers have been at the forefront in this area since adopting a fiscal rule to guide government spending decisions a decade ago. This so-called structural balance rule (SBR) incorporates fluctuations in copper prices—the main source of volatility in fiscal revenues—and was instrumental in saving a large part of the windfall during the commodity boom of 2005–08. When the country went into recession in 2009, however, the rule was essentially abandoned as authorities implemented a fiscal expansion beyond that suggested by the SBR.

While having a fiscal rule has served Chile well, there are pending questions about the appropriateness of its design. How much would welfare improve if the rule were modified to respond more to accumulated assets? Or to promote more countercyclical spending? Furthermore, since the rule is well understood and has gained legitimacy across society, it is desirable to consider improvements that do not entail major departures from its current structure. This raises the question of whether the gains from moving toward a spending policy with a higher propensity to spend out of assets when private income is low can be achieved with a rule

We thank Rómulo Chumacero for helpful comments and suggestions.

Fiscal Policy and Macroeconomic Performance, edited by Luis Felipe Céspedes and Jordi Galí, Santiago, Chile.

similar to the SBR, for example, by adding an escape clause whereby spending is expanded beyond what is prescribed for normal times in predetermined extreme circumstances.

In this paper, we explore from a normative perspective the contours of an optimal spending rule for a government that has volatile revenues from an exogenous source such as a flow from a natural resource, very much like Chile. Specifically, we analyze policies for a government with a precautionary saving motive that decides how much to transfer from volatile copper revenues to impatient agents that differ in their private incomes, which in turn are volatile and correlated with fiscal revenues. Much as in reality, the government can save abroad, has limited space for borrowing against future revenue, and has access to an imperfect technology for targeting transfers (that is, a portion of transfers leaks to richer households). Households' behavior is simple: they consume all available income.

Output is exogenous in our model, that is, fiscal multipliers are zero, so any countercyclical action reflects the desire of increasing transfers at times when household consumption is low and government spending has a higher marginal utility, rather than a Keynesian mechanism. Fiscal policy is ultimately the implementation of social insurance.

We analyze the welfare gains of an optimal rule vis-à-vis a balanced-budget rule whereby the government transfers all its revenues to households in each period. We also study the behavior of government assets and the extent to which government spending is countercyclical. We compare the optimal rule prescribed by our model with simpler rules, including the Chilean SBR, a rule that spends the permanent income from copper (à la Friedman), and linear rules similar to the SBR except that propensities to consume out of assets and structural revenues are chosen optimally. We also analyze the gains from having an escape clause.

The last global cycle made it apparent once again that government revenues in Chile are heavily influenced by copper prices. After representing less than 1 percent of GDP (or about 5 percent of total government revenues) in 1998–2003, government mining revenues increased to more than 8 percent of GDP following the rise of copper prices in 2004–08. With the subsequent decrease in commodity prices, copper revenues declined to 3 percent of GDP. Nonmining revenues, which are higher on average, have also fluctuated, but their volatility has been considerably lower. Spending decisions, on the other hand, have been guided by a predetermined central government structural surplus target (1.0 percent of GDP until 2006, 0.5 percent of GDP in 2007, and 0.0 percent in 2008). To this end, spending has been based on what is considered to be permanent revenues, stripping out cyclical revenues that include both tax revenues (influenced by the GDP cycle) and the volatile mining revenues affected by the price of copper. In principle, the rule aims to establish an acyclical fiscal behavior and the full operation of automatic stabilizers on the tax revenue side. Real government spending growth would be relatively stable and change only with innovations in trend GDP growth, changes in tax policy, and updates of what is considered the normal or reference copper price. Consequently, the overall fiscal result has varied considerably in a few years, with large savings when copper prices were high and large spending when the country went into recession (see figure 1). Government net assets increased to more than 20 percent of GDP in 2008.

Figure 1. Structural Balance Rule (SBR) and Actual Fiscal Funds

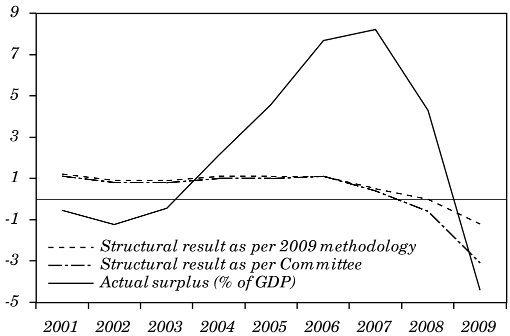

Source: Chilean Budget Office.

Fiscal policy was decisively countercyclical in 2009, when the economy entered a recession following the Lehman collapse. Real government spending increased by 18 percent (year on year), providing a fiscal impulse of 3 percent of GDP—one of the highest one-year fiscal impulses in emerging market economies during the Great Recession. Part of the fiscal reaction was in the form of targeted transfers to poor families. Unemployment increased to more than 10 percent in 2009, only slightly less than in the previous recession of 1998–99, which also followed large external shocks. Output contracted by 1.5 percent in 2009, more than the 0.8 percent drop in the previous recession. Interestingly, however, the government approval rating followed very distinct patterns: it increased significantly in 2009, largely due to perceptions of economic policies, whereas it tanked in 1998–99. This suggests that households welcome targeted fiscal policies in times of hardship.

In our model, the gains from moving from a balanced-budget rule to an optimal rule are sizeable, which indicates that the profile of fiscal spending can be quite relevant. With the baseline parameters calibrated to the Chilean economy, welfare gains from an optimal rule are equivalent to a proportional increase of copper revenues by 100 percent under a balanced-budget rule. Optimal spending displays significant countercyclicality: a fall of one standard deviation in private income leads, on average, to a rise in government transfers of 50 percent of the government's median income. The optimal rule is more countercyclical when government expenditures are less targeted, as the relative value of government transfers during recession increases in this case. Put somewhat differently, the inefficiencies of poor targeting are less costly during recessions.

Simpler rules also provide significant welfare gains. The SBR rule attains 18 percent of the gains obtained under the optimal rule; a Friedman-type rule does somewhat better, achieving 20 percent of possible gains. Gains increase substantially with linear rules where, in contrast to the Chilean SBR and Friedman-type rules, the marginal propensities to spend out of assets and wealth are chosen optimally to reflect heterogeneous households, imperfect targeting, and borrowing constraints. The results suggest a considerably lower propensity to consume out of structural copper revenues and a higher one with respect to assets, relative to the SBR. These parameters narrow the distribution of assets. The best linear rule attains 74 percent of the gains obtained under the optimal rule. Furthermore, allowing for rules that switch between two linear regimes depending on the GDP cycle further increases welfare to 83 percent of the gain under the optimal rule. As expected, the propensity to spend out of assets and structural revenue is higher in the low GDP regime. In fact, the main difference between rules with one and two linear regimes is that the former are pretty much acyclical while the latter capture the degree of countercyclical expenditure present in the optimal rule. We interpret the quasi-optimality of a regime-switching rule as the gains from having escape clauses for extreme events, which simple rules are not able to handle adequately.

The paper is organized as follows. Section 1 provides a brief literature review. Section 2 describes the model. Section 3 implements the model with Chilean data. This section describes the optimal fiscal rule, evaluates welfare gains, and analyzes the rule's behavior under different environments and shocks. Section 4 investigates whether alternative simpler rules provide useful approximations to the optimal solution, with a special focus on Chile's structural rule and variations that could help improve it. Finally, section 5 presents some concluding remarks.

1. Relation to the Literature

This paper is related to two literatures. First, it draws from works on optimal consumption with self-insurance. The starting point is the income-fluctuation problem, where a risk averse consumer receives an exogenous, stochastic income stream and maximizes the expected discounted utility, subject to an exogenous credit constraint that assumes all debts are repaid.[1] The agent has a precautionary saving motive and is impatient, as in Zeldes (1989), Deaton (1990), and Carroll (1992, 1997).[2] The model in this paper may be viewed as an income-fluctuation problem in which a planner with volatile income saves and spends to maximize the sum of expected discounted utilities of heterogeneous, impatient households with their own volatile income sources.

This paper also relates to debate on the cost of business cycles triggered by Lucas (1987).[3] We consider a government with a highly volatile source of income and compare the welfare implications of spending incomes on receipt (balanced-budget rule) with those of using a fiscal rule. Our results show that a fiscal rule aimed at stabilizing the incomes of the poor during downturns leads to considerably larger welfare gains than those obtained by Lucas.

A second type of work connected to this paper is the study of fiscal policy rules. For the most part, the applied literature focuses

1. See Schechtman (1976) for the seminal paper and Chamberlain and Wilson (2000) for a good overview.

2. As noted by Schechtman (1976), in this setting an agent with infinite marginal utility at zero consumption optimally acts as if there were liquidity constraints even if there are none.

3. See Barlevy (2004), Lucas (2003), and Yellen and Akerlof (2006) for surveys of this literature with diverging conclusions on where it stands. Also see Krusell and others (2009) for a recent contribution.

on issues of fiscal sustainability and whether having fiscal rules is, from a positive perspective, useful to that end. IMF (2009) and Kopits (2004) are good examples of this type of analysis. The former documents that fiscal rules have become more common in recent years, with almost 80 countries having rules in place in early 2009 (versus fewer than ten in 1990), and that, on average, they have been associated with improved fiscal performance and more prudent fiscal policies. The latter work compiles several case studies to analyze conditions under which rules have succeeded and concludes that political support and transparency are critical, while the extent to which a rule is legally enshrined is largely irrelevant.

One particular strand of the fiscal policy rules literature studies the challenges arising from revenues tied to nonrenewable commodities with volatile prices (for example, oil and copper). Villafuerte, López-Murphy, and Ossowski (in this volume) analyze the recent experience with fiscal policy of commodity-rich Latin American countries; they conclude that, on average, policies have been somewhat procyclical, countries that pursued more conservative fiscal policies during the boom were able to implement more aggressive countercyclical fiscal policies during the downturn, and these dimensions of fiscal policy were not linked to fiscal rules or resource funds.

Closely related work on fiscal rules from a normative perspective focuses on commodity-related revenues. A standard approach has been to apply Friedman's permanent-income hypothesis and prescribe rules that spend the annuity value of the commodity-related wealth. Segura (2006) is one of several papers based on this approach, which is attractive because of its simplicity but has several shortcomings precisely for the same reason. Among the shortcomings is that it neglects both that households have other sources of income beyond transfers and that precautionary savings can be particularly important given commodity price volatility. Engel and Valdés (2000) analyze the intergenerational distribution of an exhaustible commodity (oil, in their case) when household income is increasing over time, as well as appropriate precautionary saving given volatile prices and imperfect insurance markets. Maliszewski (2009) applies the framework to oil-producing countries and concludes that ad hoc rules perform relatively poorly. Drexler, Engel, and Valdés (2002) apply the framework to Chile and copper, noting that actual fiscal policy has been closer to the prescriptions of a model with precautionary saving than to those of a model based solely on smoothing government expenditures. The focus in their paper is the distribution of natural resource wealth across generations, not across households over the cycle as in this paper.

Finally, a number of papers study the implications of different fiscal rules for macroeconomic volatility, including the effects of the Chilean fiscal rule, through new Keynesian dynamic stochastic general equilibrium (DSGE) models. In general, these papers assume some form of non-Ricardian behavior (so that fiscal policy has nontrivial effects) through the existence of liquidity-constrained consumers (in the form of rule-of-thumb or hand-to-mouth decisions, very much like in our model). Andrés and Domenech (2006) analyze whether there is a trade-off between the sustainability of public finances and their countercyclical power, concluding that this is not the case. Kumhof and Laxton (2009) compare a balanced-budget rule with rules that embed a more active countercyclicality, including one with a structural balance. They conclude that there are high potential welfare gains from using more active rules and that in the case of commodity-driven revenues, automatic stabilizers should be allowed to operate fully (keeping spending stable). In the specific case of the Chilean fiscal rule, both Kumhof and Laxton (2010) and Batini, Levine, and Pearlman (2009) conclude that a balanced-budget rule is inferior to a structural budget rule. The first paper also concludes that a rule with more activism than the structural balance rule lowers output volatility with a minor cost in inflation volatility but considerable movements in the fiscal instrument. None of these papers deals with imperfect targeting of fiscal policy or heterogeneous agents and the income distribution, as we do in this paper.

2. Model

We analyze the optimal program of a planner that can save and spend incomes from a natural resource to maximize the sum of discounted utilities of agents representing the economy's income quintiles. An important departure from previous work is that the planner cannot target households at will, but rather is constrained by an exogenous transfer technology.

2.1 Households

Time is discrete. Total private income follows an exogenous stochastic process, Y_t^p. Income quintiles are indexed, from the poorest to the richest, by $i = 1, 2, \ldots, 5$. Each quintile is represented by one household, all of which have a subjective discount rate of $\delta > 0$. The income share of quintile i, which remains constant over time, is denoted by s_i, with $0 \leq s_1 \leq s_2 \leq \ldots \leq s_5$ and $\sum_{i=1}^{5} s_i = 1$. Households consume all their income.[4]

2.2 Planner

The planner receives an exogenous, stochastic income stream, Y_t^g, derived from a natural resource (we could extend the model to incorporate tax revenues). The planner can save at an exogenous riskless real rate, r, with $r < \delta$, so that households (and therefore the planner representing them) are impatient.

The planner faces an exogenous debt limit, B, that allows paying back the debt with probability one, which he does.[5] That is, if the planner spends $G_t \geq 0$ in period t, beginning-of-period assets evolve according to

$$A_{t+1} = (1 + r)(A_t + Y_t^g - G_t),$$

and the borrowing constraint takes the form

$$A_{t+1} \geq -B.$$

The planner's expenditures are distributed across quintiles according to an exogenous, time-invariant, targeting function, α, so that quintile i receives $\alpha_i G$ when the planner spends G, with $\alpha_i \geq 0$ and $\sum_{i=1}^{5} \alpha_i = 1$.

2.3 Dynamic Formulations

The sequential formulation for the planner's problem at time 0 is as follows:

$$\max_{G_0, G_1, \ldots} \mathrm{E}_0 \sum_{t \geq 0} (1 + \delta)^{-t} \sum_{i=1}^{5} u\left(s_i Y_t^p + \alpha_i G_t\right),$$

4. This admittedly strong assumption allows us to avoid modelling the strategic interaction between the planner and households and provides a role for fiscal rules. We relax this assumption in Engel, Neilson, and Valdés (2011).

5. That is, B is less than or equal to the planner's natural debt limit, defined as the minimum present value of income.

subject to (Y_0^p, Y_0^g) given

(Y_t^p, Y_t^g) exogenous process, $t = 1, 2, 3, \ldots$

$A_t = (1 + r)(A_{t-1} + Y_{t-1}^g - G_{t-1}), t = 1, 2, 3, \ldots$

$A_t + B \geq 0, t = 1, 2, 3, \ldots$

$G_t \geq 0, t = 0, 1, 2, \ldots.$

The problem's recursive formulation is

$$V\left(A_t, Y_t^g, Y_t^p\right) = \max_{0 \leq G_t \leq A_t + Y_t^g + (1+r)^{-1} B} \sum_{i=1}^{5} u\left(s_i Y_t^p + \alpha_i G_t\right) + \left(1+\delta\right)^{-1} \mathrm{E}_t V\left[\left(1+r\right)\left(A_t + Y_t^g - G_t\right), Y_{t+1}^g, Y_{t+1}^p\right].$$

In periods in which the solution is interior, a straightforward calculation starting from the sequential formulation yields the Euler equation:

$$\sum_i \alpha_i u'\left(s_i Y_t^p + \alpha_i G_t\right) = \frac{1+r}{1+\delta} \mathrm{E}_t \sum_i \alpha_i u'\left(s_i Y_{t+1}^p + \alpha_i G_{t+1}\right). \quad (1)$$

The planner spends resources to equalize a weighted sum of current marginal utilities with the corresponding discounted expected weighted sum of the next period's marginal utilities. The weights are given by the targeting function, in which quintiles that benefit more from government expenditures receive a higher weight. The Euler equation also shows that an increase in expected future private incomes leads to higher current spending by the planner.

In contrast to equation (1), in periods in which the borrowing constraint is binding, we have

$$\sum_i \alpha_i u'\left(s_i Y_t^p + \alpha_i G_t\right) > \frac{1+r}{1+\delta} \mathrm{E}_t \sum_i \alpha_i u'\left(s_i Y_{t+1}^p + \alpha_i G_{t+1}\right). \quad (1)$$

2.4 Perfect Targeting

One of the main departures from the literature in this paper is to allow for imperfect targeting. This motivates considering first the case with perfect targeting, which requires allowing the α_i to vary over time and will serve as a useful benchmark.

When the planner can target expenditures at will, there is a simple characterization of the distribution of government expenditures across households, conditional on the choice of G_t.[6] Expenditures are distributed across quintiles so as to equalize marginal utilities among the poorer quintiles until G_t is exhausted. Richer quintiles do not receive any transfers while the remaining households achieve a common consumption level, so that poorer quintiles receive higher transfers.

More precisely, using $\tilde{G}_k$ to denote total transfers needed to equalize total incomes of quintiles 1 through k with private income of quintile $k + 1$, a straightforward calculation shows that

$$\tilde{G}_k = \sum_{i=1}^{k} i\left(s_{i+1} - s_i\right)Y_t^p,$$

where $k = 1, 2,\ldots, 4$ and where we adopt the convention that $s_0 = 0$, $\tilde{G}_0 = 0$, and $\tilde{G}_5 = \infty$.

Since the sequence $\tilde{G}_k$ is increasing, given a level $G \geq 0$ of government expenditure there is a unique nonnegative integer k such that $\tilde{G}_k \leq G < \tilde{G}_{k+1}$. The optimal allocation of G_t across quintiles transfers resources only to quintiles 1 through $k + 1$, and it does so in a way that equalizes their total incomes. Using G_i to denote the transfer to quintile i, this means that

$$G_i = \left(s_{k+1} - s_i\right)Y_t^p + \frac{G_t - \tilde{G}_k}{k+1}.$$

It follows that finding G_t is equivalent to solving a standard incomes fluctuation problem, in which the planner's instantaneous marginal utility from government expenditures is equal to

6. Engel and Valdés (2000) derive a similar result in a model that distributes natural resource wealth across generations.

$$u'\left(s_{k+1}Y_t^p + \frac{G_t - \tilde{G}_k}{k+1}\right),$$

with k given by the piecewise constant, increasing function of G_t described above.

3. Implementation and Results

In this section, we implement the model described in section 2 using data from Chile. The trusting (or impatient) reader can skip section 3.1, which describes our parameter and functional choices, and move directly to section 3.2 on the optimal policy.

3.1 Parameter Choices

To determine the joint process of private and government revenues, we considered annual data for the 1990–2009 period. We proxied Y^p by the difference between GDP and government expenditures per capita (based on data from the Central Bank of Chile), and detrended log Y^p using a quadratic trend. The resulting stationary variable is denoted by y^p in what follows. We work with detrended Y^p to highlight the relation between cyclical fluctuations and optimal fiscal policy.

We proxied Y^g by per capita fiscal revenues derived from copper, both directly from state-owned CODELCO and indirectly via taxes on privately held copper companies, using data from the Chilean Budget Office. We denote log Y^g as y^g.

We fitted a first-order vector autoregression (VAR) to (y^p, y^g). Under the identifying assumption that current innovations to y^p have no effect on current y^g, we found no statistically significant effect of past innovations of y^p on y^g (see figure 2 for the resulting impulse response functions). For our benchmark income process, we therefore chose a specification of the form

$$y_t^p = F_0^p + F_{pp}y_{t-1}^p + F_{pg}y_{t-1}^g + \varepsilon_t^p$$

and

$$y_t^g = F_0^g + F_{gg}y_{t-1}^g + \varepsilon_t^g,$$

where only contemporaneous innovations are allowed to be correlated. Section 3.3 considers two alternative specifications, one in which past values of y^g have no effect on current y^p ($F_{pg} = 0$) and the other in which past values of y^p influence y^g.[7]

Figure 2. Impulse Responses of Government Copper Revenues and Private Income[a]

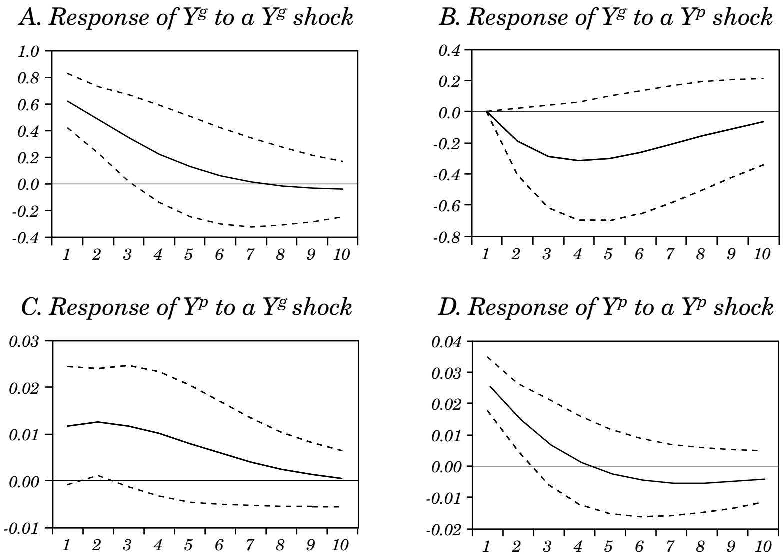

Source: Authors' computations.
a. Cholesky decomposition with Y^g the most exogenous. Dashed lines are +/– two standard deviations.

Since we are interested in fiscal rules that are relevant in coming years, we set the average value of Y_g at 2.1 percent of the average value of Y_p (which is somewhat lower than the 3 percent observed in the data) to account for the fact that Y^p was much higher toward the end of the period than at the beginning.

7. The latter could reflect, for example, a negative shock to private income that leads to a depreciation of the peso, thereby increasing revenues from copper measured in pesos. Alternatively, a negative GDP shock might cause the government to ask CODELCO to lower its investment and increase transfers to the government. As mentioned above, our VAR analysis found no statistically significant effect of past GDP shocks on current copper revenues, but the estimated coefficients are economically significant, which, given the relatively short series at hand, suggests this case may be relevant as well.

The planner's problem is solved using a Tauchen discretization for the joint distribution of (y^p, y^g). This discretization has 25 states: y^p takes five values, and there are five possible values of y^g associated with every value of y^p. Table 1 shows the probabilities of the five values of y^p and the magnitudes of the corresponding deviations from trend.

Table 1. Private Income States in the Discretization of the State Space
(in percent)

State	*Probability*	*Deviation from trend*
1	2.12	–11.9
2	22.83	–6.2
3	50.10	0.0
4	22.83	6.2
5	2.12	11.9

Source: Authors' computations.

We set the annual risk-free interest rate, r, at 5 percent and the subjective discount factor, δ, at 8 percent. A useful way to capture the notion that poor households value having smoother consumption across periods and states of nature relatively more than wealthier households is to consider an instantaneous utility function, u, that is a Stone-Geary extension of a constant-elasticity-of-intertemporal-substitution felicity function:[8]

$$u(c) = \begin{cases} \frac{1}{1-\theta}\left(c - c^*\right)^{1-\theta}, & \theta \neq 1, \\ \log\left(c - c^*\right), & \theta = 1, \end{cases} \tag{2}$$

where c^* denotes the subsistence level. We consider a coefficient of relative risk aversion, θ, of 3 in the benchmark model and set c^* at 98

8. See, for example, Deaton and Muellbauer (1980, chap. 3). An alternative route is to allow for a marginal utility of consumption that is decreasing in wealth, as in Blundell, Browning, and Meghir (1991), Attanasio and Browning (1995), Atkeson and Ogaki (1996), and Guvenen (2006). We are exploring this route in ongoing work.

percent of the income of the poorest quintile in the worst aggregate income scenario, which corresponds to an annual per capita income of approximately US$800 (the poverty line varied around US$1,200 in the period considered).

To solve the model we impose an upper bound on accumulated assets equal to average private income; this restriction is rarely binding, and our results do not change when we loosen it. We also impose a lower bound of zero on assets ($B = 0$).

Table 2 shows the values for the income share and expenditure share parameters, s_i and α_i, for each quintile. They correspond to values reported by MIDEPLAN in 2009, which are calculated using the CASEN 2009 household survey. Social expenditure targeting in Chile is considerably better than in most developing countries: Rey de Marulanda, Ugaz, and Guzmán (2006, figure 1) suggest that the typical targeting function in Latin America is close to uniform targeting, that is, to having $\alpha_i = 1/5$ for all quintiles.

Table 2. Income and Expenditure Shares: Chile, 1990–2009
(in percent)

Quintile	*Income share* $(100 \times s_i)$	*Expenditure share* $(100 \times \alpha_i)$
1	3.6	44.2
2	8.3	24.6
3	12.7	16.6
4	19.6	10.3
5	55.8	4.3

Source: MIDEPLAN and CASEN (2009).

3.2 Optimal Policy

Panel A in figure 3 shows optimal government expenditure as a function of government assets, for three values of private income. Government income is held (approximately) constant at its median value.[9] The figure plots curves for high private income (highest value

9. As described above, the discretization we consider leads to small differences in y^g across the three states considered.

in the discretization), intermediate private income (median value), and low private income (lowest value in the discretization). Both G and A are normalized by average private income (referred to as average GDP in what follows). Panel B is similar except that Y^g is held (approximately) constant at its lowest value.

Figure 3. Optimal Fiscal Spending

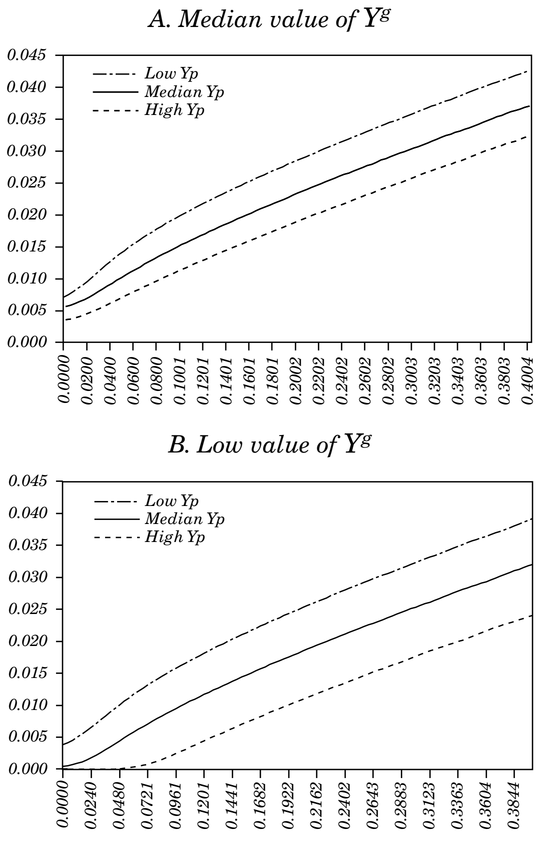

Source: Authors' computations.

Other things equal, expenditures are higher when private sector output is lower, that is, when the marginal utility of private consumption is higher. The government saves during good times to be able to spend in bad times. The expenditure functions are concave (in the regions with positive expenditure), implying a marginal propensity to spend out of assets that decreases with assets. Concavity of the expenditure function for low asset values is more pronounced during recessions (low values of Y^p), which reflects the interplay between the precautionary motive and impatience.

Comparing the two panels in figure 3 shows that government expenditures are lower when current fiscal income is lower. In fact, when fiscal revenues are low and private income is sufficiently high, there is a range of asset values in which the government finds it optimal not to spend at all.

3.2.1 Asset accumulation

Mean and median assets in steady state are equal to 38.9 and 32.9 percent of average GDP, even though assets accumulate slowly. Starting from zero, mean and median assets during the first 25 years of the rule are 13.2 and 6.1 percent of GDP, respectively. Figure 4 depicts the corresponding histogram, based on 4,000 simulations in 25 periods each (that is, 100,000 observations).

Figure 4. Distribution of Assets under Optimal Rule: First 25 Years

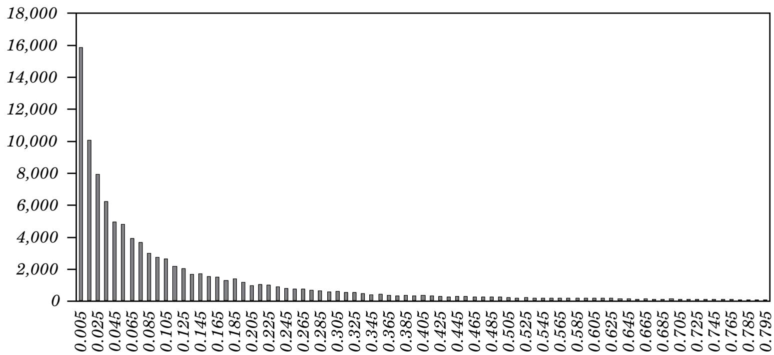

Source: Authors' computations.

3.2.2. Welfare gains

To gauge the welfare gains under the optimal rule, we quantify the associated welfare improvement with that obtained under a balanced-budget rule where the government does not incur debt or save from current income. To do this, we solve for γ in:

$$\mathrm{E}_0\left\{\sum_{t\geq 0}(1+\delta)^{-t}\sum_i u\left[s_i Y_t^p+\alpha_i(1+\gamma)Y_t^g+\alpha_i\frac{r}{1+r}A_0\right]\right\}$$
$$=\mathrm{E}_0\left[\sum_{t\geq 0}(1+\delta)^{-t}\sum_i u\left(s_i Y_t^p+\alpha_i G_t\right)\right].$$

Thus γ measures the fraction by which fiscal revenue must increase when the government spends all its income on receipt, to achieve the same level of expected welfare as under the optimal rule.[10] We obtain a value for γ of 1.001 starting from $A_0 = 0$. The welfare gain under the optimal fiscal rule is considerable.

An alternative welfare measure compares gains under the optimal rule with a scenario with no natural resource income. Using Q to denote the ratio of average government and private incomes, we solve for γ^* in:

$$\mathrm{E}_0\left\{\sum_{t\geq 0}(1+\delta)^{-t}\sum_i u\left[s_i\left(1+\gamma^* Q\right)Y_t^p+s_i\frac{r}{1+r}A_0\right]\right\}$$
$$=\mathrm{E}_0\left[\sum_{t\geq 0}(1+\delta)^{-t}\sum_i u\left(s_i Y_t^p+\alpha_i G_t\right)\right]. \tag{3}$$

The normalization constant Q is such that $\gamma^* = 1$ when $s_i = \alpha_i$ for all i and the natural resource income is equal to a constant fraction of private income, in all periods and for all quintiles ($Y_t^g = \lambda Y_t^p$). Starting with no assets, the value of γ^* is equal to 3.122 for the optimal program. Thus, even though copper revenue equals only 2.1 percent of GDP, on average, the welfare improvement it fosters

10. When $A_0 > 0$, we assume that in the balanced-budget counterfactual the government spends the annuity value from A_0.

under the optimal fiscal rule is akin to increasing private income by 6.6 percent. This happens because targeting is considerably better than having transfers proportional to quintile income and because the natural resource revenue is far from perfectly correlated with private income (correlation of 0.45). It is also possible to calculate γ^* for the balance-budget rule, by solving equation (3) with

$$G_t = Y_t^g + \frac{r}{1+r} A_0.$$

The solution is denoted by γ^*_{BB} and equal to 1.65 in our baseline, implying that welfare under a balanced-budget rule is the same as under a 3.5 percent increase in private incomes and no natural resource revenue ($3.5 \cong 1.65 \times 2.1$ percent).

3.2.3 Cyclical behavior

The macroeconomic implications of the optimal fiscal rule for the cyclical behavior of government expenditure can be captured in various ways. Obvious options are the correlation between the economic cycle (as measured by detrended Y_t^p) and either government expenditures or government savings. We would expect the latter to be procyclical and the former to be countercyclical.

Denoting government saving by S_t, we have $S_t = Y_t^g - G_t$, and a straightforward calculation shows that

$$\sigma(S_t)\rho(S_t, Y_t^p) + \sigma(G_t)\rho(G_t, Y_t^p) = \sigma(Y_t^g)\,\rho(Y_t^g, Y_t^p), \tag{4}$$

where $\rho(x_t, y_t)$ denotes the time-series correlation between x_t and y_t, while $\sigma(x_t)$ denotes the standard deviation of x_t. Equation (4) shows that procyclical government saving is equivalent to countercyclical government spending only when private and government income are uncorrelated. When the two sources of income are positively correlated—as is the case in most countries with significant revenues from natural resources, including Chile—the possibility of procyclical saving and expenditure arises. This is the case for the optimal policy in our benchmark model: the correlation between government saving and the economic cycle is 0.30, while the correlation between government spending and the cycle is 0.26. By comparison, these correlations are zero and 0.45, respectively, for a balanced-budget rule.

An alternative way to quantify the extent to which optimal spending varies with the business cycle, Y^P, is to use ordinary least squares (OLS) to estimate a linear approximation to the optimal rule of the form

$$G_t = c_0 - c_p Y_t^p + c_g Y_t^g + c_a A_t + \text{error}$$

and measure the degree of countercyclicality by

$$CCG \equiv c_p \times \frac{\sigma\left(Y^p\right)}{\text{med}\left(Y^g\right)}, \tag{5}$$

where med(x_t) denotes the median of x_t. *CCG* captures the response of government expenditures, as a fraction of median government income, associated with a decrease of one standard deviations in private income. For the benchmark model we obtain $CCG = 0.49$, which implies that government expenditure, as a fraction of median government income, increases by 49 percent, on average, when private income drops by a standard deviation.

3.3 Alternative Parametrizations

Column 1 in table 3 shows the main statistics for the benchmark model: welfare gains, compared with both a balanced-budget rule and a scenario with no natural resource income (γ and γ^*); measures of asset accumulation under the optimal rule (median accumulation during the first 25 years and in steady state); two indices for countercyclical behavior (namely, the correlation between savings and the cycle and the *CCG* measure defined in equation 5); and welfare gains under a balanced-budget rule, γ^*_{BB}. The first three and the last statistic assume initial assets equal to zero; the fourth, fifth, and six rows report steady-state values.

Columns 2 through 8 show summary statistics for the optimal rule if we modify parameters from the benchmark model that characterize household preferences, one at a time. The cost of moving to the optimal rule when the initial level of assets is low is front loaded, since the planner must accumulate assets to spend in times when the marginal utility of consumption is high. By contrast and for the same reason, the benefits of adopting a fiscal rule are back loaded. This explains why an increase in households' subjective discount factor lowers welfare gains and asset accumulation (column 2), while a decrease has the opposite effect (column 3).

Table 3. Alternative Preferences

Variable	*Benchmark* (1)	$\delta = 0.1$ (2)	$\delta = 0.06$ (3)	$\theta = 5$ (4)	$\theta = 1$ (5)	$c^* = 0$ (6)	$c^* = 50\%$ (7)	$c^* = 90\%$ (8)
γ	1.001	0.806	1.307	1.526	0.131	0.176	0.275	0.557
γ^*	3.244	2.920	3.747	1.199	4.222	6.705	6.793	5.093
Med(A_{25})	0.061	0.054	0.070	0.024	0.015	0.006	0.016	0.044
Med(A_{ss})	0.329	0.225	0.512	0.055	0.047	0.039	0.105	0.264
$\rho(S, Y^P)$	0.299	0.288	0.299	0.196	0.232	0.236	0.266	0.295
CCG	0.491	0.412	0.557	0.809	0.117	0.097	0.262	0.451
γ^*_{BB}	1.647	1.635	1.662	0.594	3.836	5.919	5.672	3.642

Source: Authors' computations.

The elasticity of intertemporal substitution for the instantaneous utility defined in equation (2) is

$$\text{EIS} \equiv -\frac{u'(c)}{cu''(c)} = \frac{c - c^*}{\theta c},$$

which is decreasing in θ and c^*. This explains why columns 4 through 8 show that the benefits of a fiscal rule are larger when households have a stronger preference for a smoother consumption over time (smaller EIS). Also, fiscal policy is more countercyclical when households, particularly those in the poorest quintile, are less able to smooth consumption over time. The countercyclical measures of fiscal policy are significant in all cases, although they are sometimes smaller than for the benchmark model.

Table 4 considers changes in the income processes. Columns 2 and 3 fit separate AR(1) processes to y^p and y^g and assume independent innovations (column 2) and correlated innovations (column 3, where the correlation is 0.40). The value of owning copper, compared with a scenario with no natural resource revenues, is larger when innovations are independent than when they are positively correlated, both under the optimal policy and under a balanced-budget policy (γ^*_{BB} of 3.81 versus 2.45; γ^* of 5.24 versus 4.12). The reason for this is that a revenue stream that is uncorrelated with private income provides more insurance than a positively correlated income source.

Column 4 of table 4 considers a first-order VAR in which past private income shocks are allowed to affect current commodity revenues (see footnote 7). Specifically, revenues from the natural resource can be expected to rise in the period following a negative innovation to private income, which allows the planner to spend more aggressively today, since there is less need to save resources for future periods. This explains why the value of the optimal policy, as measured by both γ and γ^*, is higher than in the benchmark case and the cases with standard AR specifications.

Columns 5 through 8 show that the benefits of a fiscal rule increase with the volatility of both fiscal and private income, compared with a balanced-budget rule, leading to higher values of γ. In the case of a change in the volatility of fiscal revenues, this improvement largely reflects the fact that the value of a balanced-budget rule deteriorates significantly when volatility increases (see the γ^* reported in the last

Table 4. Alternative Income Processes

		Independent	Correlated	Unrestricted	$\sigma(y^p)$		$\sigma(y^g)$		$\mu(Y^g)$	
	Benchmark	*AR*	*AR*	*VAR*	↓25%	↑25%	↓25%	↑25%	↓50%	↑50%
Variable	*(1)*	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*	*(8)*	*(9)*	*(10)*
γ	1.001	0.546	0.897	1.170	0.882	1.093	0.458	1.837	1.153	0.906
γ^*	3.244	5.245	4.211	4.409	3.152	3.379	5.416	1.792	3.542	3.039
$\text{Med}(A_{25})$	0.061	0.040	0.051	0.053	0.060	0.061	0.054	0.060	0.038	0.086
$\text{Med}(A_{ss})$	0.329	0.263	0.293	0.270	0.336	0.326	0.253	0.369	0.174	0.446
$\rho(S,Y^P)$	0.299	0.070	0.271	0.296	0.283	0.314	0.338	0.260	0.332	0.283
CCG	0.491	0.304	0.390	0.613	0.437	0.521	0.216	1.133	0.481	0.476
γ^*_{BB}	1.647	3.813	2.451	2.245	1.728	1.617	3.735	0.695	1.671	1.635

Source: Authors' computations.

row of the table): an increase in the volatility of copper revenues, which is positively correlated with private income, decreases the extent to which this income stream provides insurance.

Columns 9 and 10 consider changes in the importance of copper revenues, with a 50 percent decrease and increase, respectively. The value of the optimal program, compared with the balanced-budget rule, is larger when copper revenue is less important. The marginal benefit of additional natural resource income is smaller when overall resources are larger, as these resources are likely to be spent at times when marginal utility of additional government expenditures is lower.

Table 5 summarizes the effects of changes in the targeting technology. Welfare gains increase when Chile's targeting parameters are replaced by less focalized uniform targeting ($\alpha_i = 1/5$ for all i), while countercyclicality increases considerably. The relative social value of targeting during recessions is much higher when targeting is poor. Welfare gains also increases considerably under perfect targeting (γ^*in the second row provides the correct measure in this case).

Table 5. Alternative Targeting Technologies

Variable	*Benchmark* *(1)*	*Uniform targeting* *(2)*	*Perfect targeting* *(3)*
γ	1.001	1.297	2.941
γ^*	3.244	1.713	6.155
Med(A_{25})	0.061	0.077	0.06
Med(A_{ss})	0.329	0.209	0.329
$\rho(S, Y^P)$	0.299	0.367	0.293
CCG	0.491	0.865	0.455
γ^*_{BB}	1.647	0.759	3.634

Source: Authors' computations.

Finally, table 6 provides an alternative comparison of the three targeting technologies. It reports average expenditures for the five private income scenarios (see table 1). The first column considers a balanced-budget policy, where no effort is made to use copper income to smooth household consumption or provide precautionary saving.

The remaining columns consider the same targeting technologies as in table 5. The last row of this table shows that, as expected, total expenditures are higher when the government accumulates assets. Given the extremely high volatility of copper revenues and its positive correlation with private incomes, this results in highly procyclical government transfers, explaining the dramatic difference between columns 2–4 and column 1. Government expenditures when private income is low increase considerably (by a factor between 6 and 20, depending on the policy and low income state considered) when the government moves beyond a balanced-budget policy.

Expenditures are more countercyclical when targeting is less. For example, government transfers are at least 10 percent higher under Chile's relatively good targeting than under uniform targeting. Nonetheless, expenditures are highest, on average, when private income is highest. The reason is that copper revenues are procyclical and highly persistent, so that the wealth effect associated with high copper revenues dominates the precautionary motive.

Table 6. Average G Conditional on y^p for Alternative Targeting Technologies
(in percent)

	Targeting			
	BB	*Uniform*	*Chile*	*Perfect*
Private income level	*(1)*	*(2)*	*(3)*	*(4)*
Low	0.19	3.64	3.05	2.98
Below average	0.55	3.55	3.19	3.16
Average.	1.54	3.65	3.59	3.58
Above average	4.19	4.49	4.74	4.77
High	11.30	8.97	9.36	9.38
Overall average G (%)	2.10	3.93	3.87	3.87

Source: Authors' computations.

4. Simple Rules

In practice, fiscal rules should be simple for a number of reasons. First, it is easier to communicate the constraints imposed on public spending to elected officials and the public in general when the rule is relatively simple. This helps legitimize the rule and makes it less likely that the rule will be abandoned. Second, fiscal rules are often written into laws, and this is not easy with rules that require tabulating values to characterize how much is spent and how much is saved in a given year, as in the example plotted in figure 3. That is, to be useful, rules need to be easily replicable in terms of their calculation. Third, as in the Chilean case, the starting point is often a simple rule that has earned legitimacy among policymakers and the public, so moving to a much more complex rule may come at the cost of losing this social capital.

4.1 Rules Considered

Our starting point is a version of the Chilean structural balance rule (SBR), and the question we address is how much closer we can get to the optimal rule discussed in section 3.2 with a simple variant of the SBR.

Our version of Chile's structural balance rule is written as follows:

$$G_t = \mathcal{S}_t^G + \frac{r}{1+r} A_t, \tag{6}$$

where S_t^G is the structural government income, defined as[11]

$$\mathcal{S}_t^G = \frac{1}{10} \sum_{k=0}^{9} \mathrm{E}_t Y_{t+k}^G ,$$

where E_t denotes expectations based on information available in period t, which in our case is current and past values of both income processes. The SBR prescribes spending the sum of the current structural income,

11. We focus on copper-related revenue and continue ignoring tax revenue. In practice, every year the Finance Minister appoints a committee of experts that provides an estimate for S_t^G. See Frankel (in this volume) for a discussion of the institutional design of the rule.

equal to the best estimate for average income over the next decade, and the (long-term) interest obtained on assets saved.

The SBR is similar to the optimal spending/saving rule implied by Friedman's permanent income theory of consumption, with structural income in place of wealth. For this reason, we also consider the following Friedman-type rule:

$$G_t = \frac{r}{1+r}\left(\mathcal{W}_t^G + A_t\right), \tag{7}$$

where

$$\mathcal{W}_t^G = \sum_{k \geq 0} \left(1+r\right)^{-k} \mathrm{E}_t Y_{t+k}^G$$

denotes government wealth.

We consider the following simple variant of the SBR, which keeps the basic linear structure but frees up the values for the marginal propensities:

$$G_t = c_0 + \theta_s \mathcal{S}_t^G + \theta_a A_t. \tag{8}$$

Equation (8) defines a rule that is linear in structural income and assets, but optimizes over the corresponding coefficients.

As mentioned in the introduction, real government spending increased by 18 percent (year on year) in 2009, going beyond the increase suggested by the SBR and providing a fiscal impulse of 3 percent of GDP. Some analysts argued at the time that this increase could be justified by the fact that the SBR did not allow for a marginal propensity to spend out of assets that increased during recessions.[12] This motivates considering linear spending rules with coefficients that vary with the level of private income, such as

$$G_t = c_0 + \begin{cases} \theta_{sl} \mathcal{S}_t^G + \theta_{al} A_t, & \text{if } Y^p \text{ is low;} \\ \\ \theta_{sh} \mathcal{S}_t^G + \theta_{ah} A_t, & \text{if } Y^p \text{ is normal or high.} \end{cases} \tag{9}$$

12. See, for example, "Eduardo Engel y los vientos económicos," *La Segunda,* 24 July 2009, p. 40.

The marginal propensities are allowed to vary with the economic cycle, as captured by private income, Y^P. We consider the case in which these coefficients can take two (optimally chosen) values, depending on whether private income is low (namely, the lowest two values in table 1) or normal/high (the highest three values in table 1).

Rule (9) is a regime-switching rule with two simple linear regimes, which can be thought of as a rule with an escape clause. A simple linear rule operates most of the time (75 percent in our case), but it is abandoned in extreme circumstances, when private income (in deviation from trend) is below a certain threshold.

As with all the simple rules we study in this section, we impose the same borrowing constraints considered when deriving the optimal rule in section 2, that is, $A_t \geq 0$ and $G_t \geq 0$.[13] To estimate the parameters in models (8) and (9) we first generate 1,000 time-series for private and government income, each with 100 observations: $Y^p_{k,t}$ and $Y^g_{k,t}$, $k = 1,\dots 1{,}000$; $t = 1,\dots 100$. Next we use the Nelder-Mead simplex method to find the parameter configuration, θ, within the family of rules being considered, Θ, that maximizes $\gamma(\theta)$, defined via

$$\sum_{k=1}^{1{,}000}\sum_{t=0}^{99}(1+\delta)^{-t}\sum_{i=1}^{5}u\left\{s_i Y^p_{k,t} + \alpha_i\left[1+\gamma(\theta)\right]Y^g_{k,t} + \alpha_i \frac{r}{1+r}A_0\right\}$$

$$= \sum_{k=1}^{10{,}000}\sum_{t=0}^{99}(1+\delta)^{-t}\sum_{i=1}^{5}u\left[s_i Y^p_{k,t} + \alpha_i G\left(A_{k,t}, S^g_t, Y^p_{k,t};\theta\right)\right],$$

where $A_{k,t}$ denotes the value of assets and $G(A_{k,t}, S^g_t, Y^p_{k,t};\theta)$ optimal expenditure, both for the kth time series, under rule $\theta \in \Theta$ at time t. This determines the optimal rule, $\hat{\theta}$. To avoid overfitting, the value of γ we report for $\hat{\theta} \in \Theta$ is obtained by rerunning the above procedure with 4,000 series of newly generated income series of length 100 each.

4.2 Results

Table 7 presents the summary statistics for the simple rules considered in this section. The SBR and the Friedman-type rule attain 18 and 20 percent of the welfare gain obtained under the

13. Thus, for example, the rule in equation (8) actually has $G_t = \max(0, c_0 + \theta_s S_t^G + \theta_a A_t)$.

optimal rule, respectively. These rules tend to underaccumulate assets when compared with the optimal rule, and, not surprisingly, both of them vary very little, if at all, with the economic cycle.[14]

Table 7. Simple Rules

Rule	*Welfare gain* γ ($A_0 = 0$)	*Steady-state* *Median assets*	*CCG*
Benchmark	1.001	0.329	0.491
Chile's SBR	0.180	0.095	–0.159
Friedman	0.205	0.161	–0.001
Linear rule (8)	0.743	0.160	0.092
Rule with exit clause (9)	0.830	0.154	0.454

Source: Authors' computations.

An SBR-type rule, in which the marginal propensities to spend out of current government income and assets are chosen optimally, leads to higher welfare, with approximately 74 percent of the gain under the optimal rule. Table 8 reports the estimated marginal propensities to consume out of assets in this case, showing that the improvement in performance is achieved by more than doubling the propensity to spend out of assets and reducing by more than two-thirds the propensity to spend out of structural income. This suggests that the SBR is too responsive to changes in structural income and responds too little to changes in assets. This insight is robust across specifications: the median value for the marginal propensity to spend out of assets across the 19 models considered in tables 3, 4, and 5 is 0.117, with an interquartile range of 0.025. Similarly, the median value for the propensity to spend out of structural revenue is 0.335, with an interquartile range of 0.293 (the range of values goes from 0.116 to 0.747).

The regime-switching rule achieves a significant welfare gain, attaining 83 percent of the gains obtained under the optimal rule (with a γ of 0.830 versus 1.001). Both rules accumulate considerably

14. In fact, the SBR is somewhat procyclical, reflecting the fact that structural revenue is procyclical and that the linear term in assets is not important enough to undo this effect.

fewer assets than the optimal rule. More important, the rule with an exit clause achieves a degree of countercyclicality similar to that of the optimal rule, while the optimal linear rule does not.

Table 8 also shows that the propensities to spend out of the government's assets under the rule with an exit clause are considerably larger during recessions than under the linear rule, where these propensities are chosen optimally but are not allowed to vary over the cycle. By contrast, the propensities to spend during expansions are similar under both rules where this propensity is chosen optimally. With regard to the propensity to spend out of structural income, rule (9) has a higher propensity during recessions than rule (8), but a lower propensity during normal times or expansions. A linear rule has a hard time capturing the countercyclical behavior of the optimal rule, while a rule with an exit clause can capture this feature with a marginal propensity that is higher when income is low.

Table 8. Simple Rules and Marginal Propensities to Spend

Rule	*A*	S^G	*Constant*
Chile's SBR:	0.048	1.000	—
Linear rule (8)	0.118	0.290	–0.0006
Rule with exit clause (9)			
Y^P low:	0.164	0.467	–0.0023
Y^P normal or high:	0.120	0.261	–0.0023

Source: Authors' computations.

The above insight can be applied to gauge how much government expenditures should have increased when the economy went into recession in 2009. The rule with an escape clause suggests an increase of almost one percentage point of GDP higher than the increase implied by the linear rule when accumulated assets are 20 percent of GDP, which was the level of the Chilean government's net assets going into 2009. Similarly, assuming structural government revenue was at its average value of 2.1 percent, moving to the linear rule with an escape clause leads to additional expenditures of approximately 0.4 percent of GDP. The combined effect is an increase of 1.4 percent of GDP beyond that suggested by the rule in normal times, a meaningful fiscal expansion.

Summing up, a simple linear rule with an exit clause (which leads to a different, equally simple, linear rule) does a remarkably good job at capturing the nonlinearities present in the optimal policy. Furthermore, this rule leads to lower asset accumulation and can be explained as a straightforward generalization of the SBR. Both these factors should enhance its political viability.

5. Conclusion

We have explored the qualitative and quantitative implications of different ways to conduct fiscal policy, that is, the decision of how much to spend out of government income, in a framework in which fiscal expenditure has nontrivial effects because households are hand-to-mouth consumers and both household and government incomes face unpredictable shocks. Government income is particularly volatile, as it depends on the price of a primary commodity.

The basic intuition guiding government expenditures is straightforward: the authorities seek to help the private sector smooth consumption by combining a precautionary motive with the smoothing of transitory income shocks (à la Friedman). However, the government does not only consider its own revenue and assets when deciding how much to spend, but also looks at how the private sector is doing, spending more when the private sector's income is low. Furthermore, because there is income heterogeneity across households, and the government has only a limited ability to transfer income to the poor, the government faces a nontrivial tradeoff when implementing its spending rule: imperfect targeting increases the level of expenditure needed to achieve a given level of consumption for the poorest households, which in turn makes the optimal policy more countercyclical than if targeting were perfect. It follows that better targeting leads to less countercyclical government spending, implying that countries that have less capacity to target transfers should run a more countercyclical rule.

The application of our model to Chile, using plausible parameters for income fluctuations and correlations, the household income distribution, and the targeting technology, allows us to quantify the welfare benefits of different alternatives for conducting fiscal policy, from a (complex) optimal policy function to simple linear rules, including the Chilean structural balance rule (SBR). In comparison with a balanced-budget rule, the optimal rule improves welfare by the equivalent of a 100 percent increase of government copper revenue per year under our baseline calibration, which includes positive effects from copper prices to private sector income. The optimal policy involves significant expected asset accumulation as a buffer stock, equivalent to around 33 percent of GDP in our baseline, although it takes many years to reach large values. More important, the optimal policy implies a considerable degree of countercyclicality: a fall in private income of one standard deviation translates, on average, into a 50 percent rise in government transfers relative to median government income. In certain states (characterized by high private income, low copper revenues, and low assets), the optimal policy is to save all current income and cut transfers to zero.

The SBR used in Chile over the past decade and a Friedman-type rule attain meaningful welfare gains of around 20 percent of those achieved by the optimal rule. On average, both simple rules accumulate fewer assets than the optimal policy and are close to acyclical. Optimizing the marginal propensities to spend out of assets and structural government income for an SBR-type rule results in a propensity to spend out structural or permanent copper revenues that is much lower than one, together with a propensity to spend out of assets that is much higher than the annuity value. This rule yields considerable additional gains, attaining a surprising 74 percent of gains obtained under the optimum. The result that the Chilean rule tends to spend too much out of copper and too little out of assets is robust across parameter specifications.

Finally, motivated by the quantitative importance of the optimal rule's countercyclical behavior, we also explored the gains from a regime-switching rule with two linear rules, which allows for higher spending when household income is particularly low (private sector in recession). This higher spending in certain states of nature obviously needs higher savings in normal times. The welfare gain in this case is a surprising 83 percent of the optimum. The policy implication is that there would be substantial benefits from adding an escape clause to the Chilean SBR for recessions, when countercyclical spending is valued most, increasing the propensities to spend out of assets and structural income, even though the latter remains below one. The fact that the SBR was effectively abandoned in 2009 may not be coincidental, as it allowed the rule to provide social insurance.

References

Andrés, J. and R. Domenech. 2006. "Fiscal Rules and Macroeconomic Stability." *Hacienda Pública Española* 176(1): 9–42.

Atkeson, A. and M. Ogaki. 1996. "Wealth-Varying Intertemporal Elasticities of Substitution: Evidence from Panel and Aggregate Data." *Journal of Monetary Economics* 38(3): 507–34.

Attanasio, O. and M. Browning. 1995. "Consumption over the Life Cycle and over the Business Cycle." *American Economic Review* 85(5): 1118–37.

Barlevy, G. 2004. "The Cost of Business Cycles and the Benefits of Stabilization: A Survey." Working Paper 10926. Cambridge, Mass.: National Bureau of Economic Research.

Batini, N., P. Levine, and J. Pearlman. 2009. "Monetary Policy Rules in an Emerging Small Open Economy." Working Paper 09/22. Washington: International Monetary Fund.

Blundell, R., M. Browning, and C. Meghir. 1991. "Consumer Demand and the Life-Cycle Allocation of Household Expenditures." *Review of Economic Studies* 61(1): 57–80.

Carroll, C. 1992. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity* 2: 61–135.

———. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112(1): 1–55.

Chamberlain, G. and R. Wilson. 2000. "Optimal Intertemporal Consumption under Uncertainty." *Review of Economic Dynamics* 3(3): 365–95.

Deaton, A. 1990. "Saving and Liquidity Constraints." *Econometrica* 59(5): 1221–48.

Deaton, A. and J. Muellbauer. 1980. *Economics and Consumer Behavior*. Cambridge University Press.

Drexler, A., E. Engel, and R. Valdés. 2002. "Copper Uncertainty and the Optimal Fiscal Strategy in Chile." In, *Análisis Empírico del Ahorro en Chile*, edited by F. Morandé and R. Vergara. Santiago: Central Bank of Chile.

Engel, E., C. Neilson, and R. Valdés. 2011. "Fiscal Rules as Social Policy." Yale University.

Engel, E. and R. Valdés. 2000. "Optimal Fiscal Strategy for Oil-Exporting Countries." Working Paper 00/118. Washington: International Monetary Fund.

Guvenen, F. 2006. "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective." *Journal of Monetary Economics* 53(7): 1451–72.

IMF (International Monetary Fund). 2009. "Fiscal Rules: Anchoring Expectations for Sustainable Public Finances." Washington: IMF, Fiscal Affairs Department.

Kopits G., ed. 2004. *Rules Based Fiscal Policy in Emerging Markets*. New York: Palgrave MacMillan.

Krusell, P., T. Mukoyama, A. Sahin, and A. Smith. 2009. "Revisiting the Welfare Effects of Eliminating Business Cycles." *Review of Economic Dynamics* 12: 393–404.

Kumhof, M. and D. Laxton. 2009. "Simple, Implementable Fiscal Policy Rules." Working Paper 09/76. Washington: International Monetary Fund.

————. 2010. "Un modelo para evaluar la regla de superávit fiscal estructural de Chile." *Economa Chilena* 13(3): 5–32.

Lucas, R. 1987. *Models of Business Cycles*. Oxford: Basil Blackwell.

————. 2003. "Macroeconomic Priorities." *American Economic Review* 93(1): 1–14.

Maliszewski, W. 2009. "Fiscal Policy Rules for Oil-Producing Countries: A Welfare-Based Assessment." Working Paper 09/126. Washington: International Monetary Fund.

Rey de Marulanda, N., J. Ugaz, and J. Guzmán. 2006. "La orientación del gasto social en América Latina." Working Paper I-64. Washington: Inter-American Development Bank.

Schechtman, J. 1976. "An Income Fluctuation Problem." *Journal of Economic Theory* 12(2): 218–41.

Segura, A. 2006. "Management of Oil Wealth under the Permanent Income Hypothesis: The Case of São Tomé and Príncipe." Working Paper 06/183. Washington: International Monetary Fund.

Yellen, J. and G. Akerlof. 2006. "Stabilization Policy: A Reconsideration." *Economic Inquiry* 44(1): 1–22.

Zeldes, S. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97(2): 305–46.