# Transaction costs and consumption

Geng Li [1]

*Federal Reserve Board, USA*

## ARTICLE INFO

## ABSTRACT

The rational expectations permanent income hypothesis (RE-PIH) fails to explain several well documented features of consumption behavior. First, the estimated marginal propensity to consume (MPC) for unanticipated transitory income shocks is often much higher than what the theory warrants. Second, the estimated MPC is typically much bigger for small shocks of this type than for large shocks. Third, consumption is often smoothed against large anticipated future income changes but not always against small changes. This paper argues that these findings can be reconciled within a RE-PIH model that includes a cash-in-advance constraint and an assumption that the agent is required to pay a fixed transaction cost to transfer wealth between cash and assets. Key results of the paper include first, the agent follows an $s-S$ rule with respect to cash holdings when he makes wealth-transfer decisions; second, the MPC within the no-transfer band is higher than that out of the band, and can be as high as exactly equal to one; and third, the agent smoothes consumption in response to news of large future income changes but not necessarily to small ones.

Published by Elsevier B.V.

## 1. Introduction

The rational expectations-permanent income hypothesis (RE-PIH) makes two important predictions about how consumption should respond to income changes. First, under complete markets and perfect foresight assumptions, the magnitude of consumption changes in response to unanticipated transitory income shocks, which are often referred to as "windfall" income shocks, should be equal to the annuity value of the shock. Second, in response to anticipated income changes, all of the adjustments in consumption should have been carried out upon the arrival of the news about future income changes, and no further adjustments are needed.

Testing the validity of the RE-PIH has been an important part of economic research since the advent of the theory. After several decades of research, the literature has yielded mixed findings. Among the early contributions of estimating the marginal propensity to consume (MPC) out of windfall income shocks, Bodkin (1959) finds the MPC can be as high as 0.9, a level well above any reasonable annuity value. In contrast, Kreinin (1961) finds that the MPC out of a windfall income shock is around 0.15, which is much closer to what the RE-PIH predicts. Bodkin's experiment involved small income shocks, whereas Kreinin's involved much larger ones. These findings suggest that the MPC decreases as the size of the income shock becomes larger, a relationship that cannot be explained by the standard RE-PIH. Much of the more recent work has been focusing on testing the validity of the Euler equation for households with anticipated income changes. Researchers typically report that consumption does respond to anticipated income changes, a phenomenon typically referred to as

excess sensitivity of consumption, but only when the changes are relatively small (Parker, 1999; Souleles, 1999, 2002). Little excess sensitivity has been detected when the anticipated income changes are large (Souleles, 2000; Hsieh, 2003).

Various explanations have been offered for an exceedingly large estimate of MPC and for the excess sensitivity. However, explanations for the different response to large and small income shocks are lacking. This paper introduces a cash-in-advance model that predicts (1) small windfall income shocks induce a large MPC, (2) the MPC decreases as the shocks become larger, and (3) consumption is more likely to exhibit excess sensitivity in response to anticipated income changes when the changes are small.

In this model, a consumer can hold cash and illiquid assets (hereafter assets). The nominal interest rate is zero on cash and positive on assets. Only cash can be used to purchase consumption goods. The consumer has to pay a transaction cost to transfer wealth between cash and assets and the transaction cost is assumed to be a fixed fee that does not depend on the amount of the transaction. Transaction costs have long been studied and recognized as an important source of financial market frictions. For example, Vissing-Jørgensen (2003) estimates that a per-period participation cost of merely 50 dollars can explain a substantial portion of the lack of participation in the equity market. Odean (1999), studying the data from a large brokerage firm, finds that even small transaction costs can keep optimal stock trading frequency low. Moreover, Barber and Odean (2000) calculate the optimal trading frequency for reasonable levels of transaction costs, and report that the number of trades per year should be within the range of 0.17–0.5 when the transaction cost varies from 0.01% to 0.1% of the portfolio value. Mechanically, the transaction costs introduced here is similar to Alvarez et al. (2002). In the current model, the consumer will choose to pay the transaction cost only if income shocks have made the consumer's cash balance too low or too high relative to her unconstrained optimal consumption level. A cash balance that is too low will severely constrain consumption and will cause the consumer to sell assets to finance her current consumption. Conversely, a balance of (non-interest-bearing) cash that is too high will cause the consumer to transfer most of the extra cash into the assets account.

Conditional on the expected future labor income and current asset holdings, there is a band of cash holdings, commonly referred to in the literature as an $s$–$S$ band, within which the consumer will choose not to transfer wealth. The $s$–$S$ band is divided into two pieces and optimal consumption—as a function of cash balance holdings—is not smooth between the pieces. The consumption function is not even continuous.[2] At the lower piece of the band, the consumer has a cash holding level that is lower than the unconstrained optimal consumption level. Hence, within this region of cash holdings, the MPC out of one additional dollar will be exactly equal to one. When the cash balance has increased to such a level that the consumer is no longer cash constrained (the right piece of the band), the MPC is still higher than that out of the band. The reason is that if the consumer does not spend all of the cash in hand, the consumer will not receive any return on the cash saved, and may even suffer inflation loss. Therefore, by holding cash in hand, current consumption will be higher, and the profile of consumption growth will be flatter. Overall, the average MPC within the band is higher than that out of the band. I will show that the probability of being within the band is quite high, and that only a large income shock will push the consumer out of the band. Therefore the model predicts that a small windfall income shock may be associated with a very high MPC, and that the MPC decreases as the size of income shocks becomes bigger.

Turning to how consumption responds to anticipated income changes, if the consumer knows that income will be higher next period but the change is not too large, a cash-constrained consumer will not pay the transaction cost to sell assets in order to increase current consumption. Only when the future income change is sufficiently large, it becomes optimal for the consumer to pay the costs and change her current consumption. As a result, an econometrician who tests excess sensitivity will find that the Euler equation is more likely to hold when the anticipated income changes are large, but less so when the changes are small.

Finally, what is noteworthy is that the model implies a precautionary demand for cash. When the consumer transfers wealth between cash and assets, he will typically choose to hold some cash into the next period even though cash does not accrue any interest. By choosing the optimum amount of cash to hold over periods, the consumer can minimize the probability of having to pay the transaction costs again and to limit the probability of being cash constrained in the future.

The paper will proceed as follows: Section 2 reviews the studies of the MPC estimations and excess sensitivities. Section 3 sets up the model and characterize the solution. Section 4 solves and simulates the model numerically, and presents the main results. A series of sensitivity tests are also conducted. Section 5 provides some concluding remarks and sets up a future research agenda.

## 2. Evidence on the MPC and excess sensitivities

I will first review the evidence that when the income shocks are small, the MPC is much larger than what the RE-PIH predicts, and the evidence that the MPC tends to decrease as the size of the windfall income shocks become bigger. Next I will review the studies of consumption excess sensitivity in response to anticipated income changes. The literature suggests that the excess sensitivities are more likely to be detected when income changes are small, but not when they are

---

[2] Conventionally, the lower case $s$ refers to the lower threshold of the band and the upper case $S$ refers to the upper threshold of the band. Throughout this paper, I will abuse the terminology slightly by referring to the optimal consumption policy as the consumption function. More rigorously speaking, it should be referred to as a correspondence because the policy involves one-to-many type of mapping at the $s$–$S$ thresholds.