Economic Quarterly—Volume 101, Number 3—Third Quarter 2015—Pages 193–223

Understanding Living Wills

Arantxa Jarque and Kartik Athreya

During the recent financial crisis of 2007–08, several large financial institutions came close to failing. This led to a number of publicly supported rescues and other interventions involving taxpayer money.[1] In almost all of these instances, not intervening to lessen the impact of the failures on the market seemed (to some policymakers, at least) too costly. In other words, the crisis put into use the safety net for financial institutions. Ever since, fixing the so-called "too big to fail" (TBTF) problem has been a priority for policymakers.

The TBTF problem arises when a large financial institution is in financial distress: Policymakers are not generally able to commit not to rescue it from failing, mainly because of the fear of a sizeable disruption for financial markets and the economy as a whole if such a firm fails. This "ex-post" intervention of policymakers to prevent the failure, which effectively allows creditors of the firm in distress to avoid losses on their loans, implies perverse incentives for all large financial firms "ex ante": Because creditors anticipate no losses even in the event of failure, they do not make the price of their debt reflect the level of risk taken by the financial institutions. This may lead to excessive risk-taking by the firms, which in turn will mean more frequent failures, as well as more redistribution in the form of bailouts financed by taxpayers.[2]

In this article we will study how the requirement for large financial institutions to file resolution plans, or "living wills," with their regulators may help mitigate this commitment problem and, while doing

We thank the referees, Joseph Johnson, John Walter, and Zhu Wang, as well as the editor, John Weinberg, for very valuable comments. The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: Arantxa.Jarque@rich.frb.org.

[1] See White and Yorulmazer (2014) for a detailed list of these interventions.

[2] Grochulski (2011) presents a clear exposition of this commitment problem and shows how the lack of commitment may decrease welfare.

so, decrease both the frequency of failures and their negative consequences for the economy when they do happen. A living will (LW) is a document that describes how a firm would be wound down through an unassisted bankruptcy procedure in the event of financial distress in an orderly fashion and with minimal impact to the rest of the economy. Living wills are a new requirement put in place as part of the 2010 Wall Street Reform Act, also known as the Dodd-Frank Act (DFA). The DFA was crafted with the objective of preventing financial crises like the recent one from happening in the future. As part of a wide reform of financial firm regulation, DFA prescribed a range of both new and strengthened requirements and procedures that added to the portfolio of tools with which the financial firms' supervisors work to ensure a strong financial system. Two prominent examples of existing tools that were reinforced in DFA are capital and liquidity requirements. Two important examples of new tools, which we will analyze in this article, are the requirement for systemically important financial institutions (SIFIs) to file living wills annually with their regulators and the provisions for the Orderly Liquidation Authority (OLA).

The OLA provisions, described in Title II of the Act, authorize the Federal Deposit Insurance Corporation (FDIC) to manage the winding down of certain troubled financial firms. The possibility of resolution under OLA was created by DFA as an alternative to bankruptcy, in recognition of the difficulties that may arise when using bankruptcy to resolve these very large, complex, and interconnected SIFIs. Before 2010, if a financial firm was deemed insolvent or undercapitalized and was not able to attract new capital, negotiate a bail-in by its creditors, or find a buyer in the market, it had to resort to bankruptcy. If debtor-in-possession (DIP) financing for a reorganization under Chapter 11 was not available in the market, the firm was forced into liquidation under Chapter 7. If policymakers viewed bankruptcy as too costly an alternative for society, they had options to provide public support through a purchase and acquisition of selected assets of the troubled institution (usually mediated by the FDIC if they involved depository institutions, possibly with explicit assistance in the form of asset guarantees), or a taxpayer-funded bailout that injected capital or guaranteed a loan with favorable terms. OLA constitutes a new alternative in which the firm does get reorganized and liquidated, but in a more orderly and efficient manner than through bankruptcy.

Although the details of resolution through OLA are still not clear (it has never been used so far), it has been pointed out that this alternative may be convenient in times of aggregate financial distress: It allows the FDIC to borrow from the Orderly Liquidation Fund (OLF), a dedicated account at the Treasury, at low interest rates to finance the operations of the firm in distress for at least some time. The availability of these cheap funds is likely to increase the liquidation value of the firm and possibly decrease the disruptions to the market, even in situations in which otherwise necessary DIP financing would not be available. If the liquidation of the distressed firm does not provide enough resources to repay the loans from the Treasury, DFA gives authority to charge fees on the solvent SIFIs to cover the difference, so no taxpayer money is used in OLA. However, it has also been pointed out that the availability of interim funding may benefit creditors that would otherwise get hurt from a sudden liquidation, hence leaving at least some of the perverse incentives of the TBTF problem in place.[3]

Despite the creation of the OLA, DFA still establishes bankruptcy as the preferred option for resolving a SIFI that is in financial difficulties. In order to make bankruptcy a more viable and orderly alternative, DFA requires firms designated as SIFIs to file an LW annually. Resolution through bankruptcy will be more orderly, for example, if it is easy to sell subsidiaries that are in good financial health to interested third parties. This is easiest when legal hurdles are minimal and these subsidiaries do not strongly depend on services (such as IT support) provided by other parts of the firm. As another example, resolution is easier when the failing firm has access to interim financing to keep its core operations working, which adds value to the firm. These examples suggest that a good LW should, among other things, describe the complementarities between assets and economies of scope across subsidiaries and provide a clear description of financing needs. This information would be helpful in maximizing the value of the company in bankruptcy.

Regulators review these LWs and require them to be useful and realistic. Moreover, if the plan for resolution makes apparent that certain characteristics of the firm complicate its liquidation, making the plan for liquidation "noncredible," regulators can require changes to those characteristics.

Living wills are a new tool, and regulators are still in the initial stages of implementing this requirement. Over the last few years, supervisors have been learning together with the firms about the key information that needs to be included in these documents. In this article, rather than providing a detailed description of the provisions in DFA relating to LWs and resolution, we want to lay down a *framework* that will help us understand LWs. Our objective is to study the potential benefits that LWs could bring to the regulation of financial firms,

[3] See Pellerin and Walter (2012) and Jarque and Price (2015).

and the most useful ways in which the recently installed LW review process should evolve.

In our analysis, we will emphasize the two channels through which LWs should be useful. The first channel is their annual review process, which takes place in the ex-ante world. During this review process regulators are allowed to demand changes in the way that firms are conducting business (such as their size, or the number and level of interconnectedness of their subsidiaries) if they assess that these changes would make their potential resolution less disruptive for the economy. The second channel is through their role as "road maps" for resolution authorities (bankruptcy judges but also regulators if OLA is invoked) in the event of failure. That is, in the ex-post world, LWs indicate the most efficient way to resolve the firm with minimal impact on the market.

We will illustrate in the context of a simple model of the TBTF problem how regulating living wills (rejecting noncredible plans and, importantly, mandating changes that make them realistic) may change the ex-ante versus ex-post tension that leads to the TBTF problem, and hence change the severity of the moral hazard problem. We will ask the following questions: What are the properties of living wills that make them most useful as a commitment device and improve ex-ante welfare? Under what conditions are they more likely to bring about this improvement? What are the potential costs that regulators should consider?

Our work here complements recent work in DeYoung, Kowalik, and Reidhill (2013) and in White and Yorulmazer (2014). These articles also explicitly consider how different alternatives (or "technologies") for resolution affect welfare. White and Yorulmazer (2014) use a simple static model to present a review of the different interventions during the 2007–08 financial crisis. DeYoung, Kowalik, and Reidhill (2013) focus instead on the dynamic properties of the too-big-to-fail problem. They highlight that, as regulators get better at resolution, they can let large firms fail at less cost to society (i.e., they are willing to implement harsher punishments to these firms in equilibrium), which translates into less risk taken by firms, and hence less failures, being sustained in a Markov equilibrium of the repeated game.

1. FRAMEWORK

The time inconsistency problem that underlines the TBTF problem is best described by looking at the diagram in Figure 1. The diagram describes the three-period game between three players: (1) a financial firm that maximizes the expected profits of its shareholders, (2) the creditors of the firm who set the interest rate on their loans to the firm as to equate the expected return of debt and risk-free bonds, and (3) a benevolent policymaker, or "planner," who maximizes the welfare of society (i.e., the joint payoffs to the shareholders, the creditors, and the rest of society). The planner has a large budget funded by tax revenues to use in potential bailouts, as well as funds in the OLF from fees collected from financial institutions that can be used by the OLA to provide funding to wind down SIFIs in distress.

The characteristics of a firm will be summarized in a vector X, partitioned into a subset of characteristics ω over which the policymaker has control, and a subset x that is chosen freely by the firm:

$$X = (\omega, x) .$$

Timing and Strategies

In period 0, the planner moves first and sets regulation, which will determine the constraints on a subset ω of the vector of characteristics X of the firm. This regulation includes capital and liquidity requirements, and the obligation to file LWs that meet the planner's standards. Choices of the firm such as size and complexity will only be in ω if LWs are regulated; that is, we model the increased regulatory powers given by the DFA with the LW review process as an expansion of the choices $\omega \subset X$ over which supervisors have control.

In period 1, the firm chooses a subset of its characteristics, x, given that creditors choose a price R for loans that makes them indifferent between lending to the firm or buying riskless bonds, which pay an interest $\bar{R}_1$.

In period 2, nature determines the realization of an economic shock, $\theta \in \Theta$, according to the density function h, and a political shock, $\varepsilon \in \epsilon$, according to the density function g. The economic shock realization contains two elements, $\theta = \left(\theta^i, \theta^a\right)$, where θ^i represents the idiosyncratic state of the individual firm, which affects the value and/or liquidity of its assets, and θ^a represents the aggregate state of the economy, which affects the cost of funding that the firm will face in case of distress in period 2, $\bar{R}_2(\theta^a)$. The political shock, $\varepsilon = (\varepsilon_B, \varepsilon_{OLA})$, summarizes society's utility cost of providing bailouts with taxpayer money, ε_B, and with OLF funds, ε_{OLA}. The first may be influenced by factors such as the level of disagreement of voters with the transfer of taxpayer money to banks, which may partly reflect the "type" or political views of the policymaker, as well as the opportunity cost of those funds, which

Figure 1 Game Tree

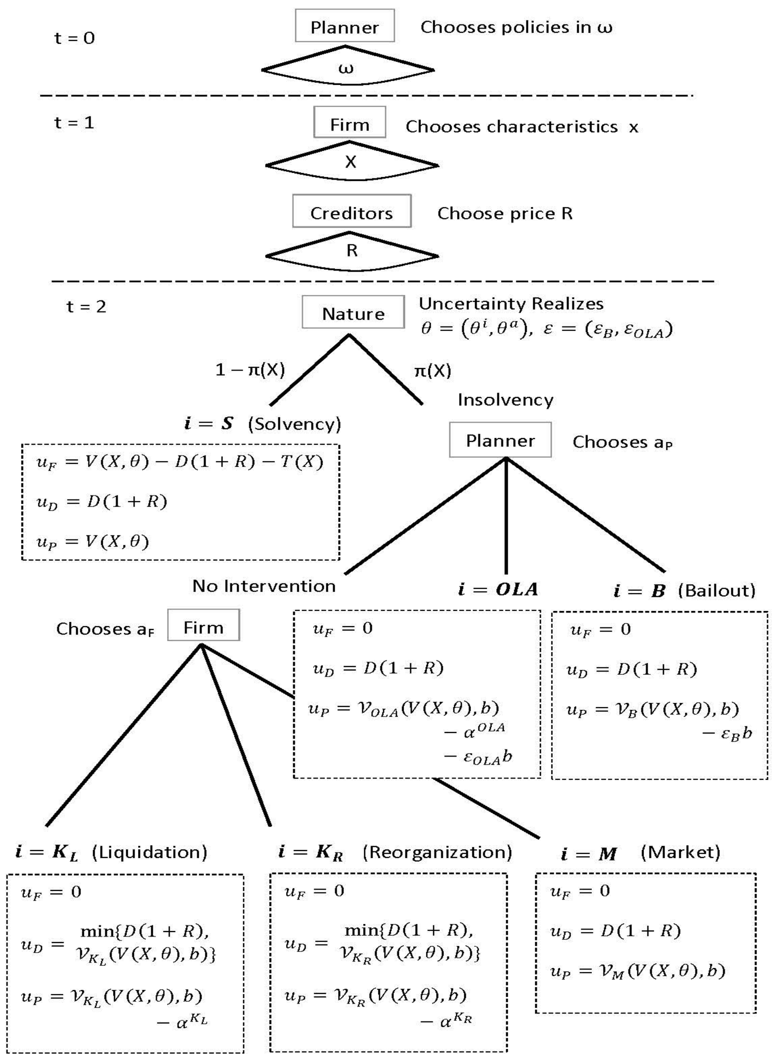

may depend on the state of public finances.[4] The second one may be influenced by factors such as the intensity of lobbying done by the banking sector against fees levied to fund the OLF, or by the opportunity cost of those funds if used to produce in the financial sector. Denoting the amount of funds devoted to help a firm in distress as b, the cost to society of providing a bailout with public funds will be $\varepsilon_B b$, while the cost of injecting funds in an OLA resolution will be $\varepsilon_{OLA} b$.

Given a realization of θ, the choices that the firm previously made, as contained in X, will determine the profit the firm makes, and hence whether it remains solvent. Denoting the gross value of the firm by $V(X,\theta)$, the amount of taxes owed by $T(X)$, the amount of debt borrowed by D, and the interest on this debt by R, we can define the set of states in which a firm with characteristics X fails, $\Theta^f(X)$, as

$$\Theta^f(X) = \{\theta \in \Theta : V(X,\theta) - T(X) - D(1+R) < 0\},$$

that is, the states in which there is not enough profit to cover tax obligations and repay creditors. The set of states in which the firm is solvent, $\Theta^s(X)$, is simply the complement of Θ^f. Using the density function over states of nature, we can calculate the probability of failure as a function of the characteristics of the firm, X :

$$\pi(X) = \int_{\Theta^f(X)} h(\theta)\, d\theta.$$

In the event that the firm is insolvent, both the planner and the firm may make choices about its resolution. First, after observing a realization θ in $\Theta^f(X)$, and the realization of ε, the planner decides whether or not to intervene. If he intervenes, he chooses whether to resolve the firm through OLA or to bail it out. If OLA is chosen, the wind down of the firm can be financed using a transfer b funded by fees levied on other financial institutions through the OLF. If a bailout is chosen instead, any funding of operations or transfers will come from taxpayer money. Second, and only if the planner chooses not to intervene, the firm chooses whether to file for liquidation or reorganization under bankruptcy, or to “go to the market.” Without modeling a market for troubled financial firms explicitly, with this last

[4] This shock could also be interpreted as a reduced-form summary or future utility cost of providing a bailout today. In particular, it could correspond to the net present expected value of the cost of future bailouts in a repeated game where the type of the policymaker has a persistent unobserved component and the market is trying to learn about it in order to correctly anticipate the policymaker’s decisions in the future. This repeated game is complicated to analyze, so we do not explicitly study it here. For a simple model of this commitment problem without learning in a repeated game, see DeYoung, Kowalik, and Reidhill (2013).

option we capture the possibility that the firm can find an interested third party to purchase its assets or bargain a merge.

These available strategies imply six final nodes of the game. One is solvency, which we denote by $i = S$. The other five nodes represent the possible outcomes after the firm is insolvent: market resolution $(i = M)$, bailout $(i = B)$, bankruptcy liquidation $(i = K_L)$, bankruptcy reorganization $(i = K_R)$, and resolution through OLA $(i = OLA)$.

The strategies of the three players in the game can be summarized as follows:

- The planner's strategy consists of a set of regulations, ω, as well as an action contingent on the realization of uncertainty, $a_P(X, \theta, \varepsilon)$: the decision to not intervene $(a_P = NI)$, intervene by providing a bailout $(a_P = B)$, or intervene by triggering OLA $(a_P = OLA)$. We denote the strategy as $\sigma_P = (\omega, a_P(X, \theta, \varepsilon))$;
- The firm's strategy consists of a set of firm characteristics that are not controlled by supervisory requirements, x, and a contingent decision of whether to look for a market buyer ($a_F = M$), file for liquidation through bankruptcy ($a_F = K_L$), or file for reorganization through bankruptcy ($a_F = K_R$). We denote this strategy as $\sigma_F = (x, a_F(X, \theta))$;
- The creditors's strategy is an interest rate: $\sigma_D = R(X)$.

Payoffs

The six end nodes of the game in Figure 1 contain the payoffs to the three players, contingent on the node i as well as X, θ, and ε. We denote them as $u_i^F(X, \theta)$ for the firm, $u_i^D(X, \theta)$ for the creditors, and $u_i^P(X, \theta, \varepsilon)$ for the planner.

In states of insolvency, $\Theta^f(X)$, the net value of the firm, $V(X, \theta) - T(X)$, is not enough to repay the debt in full. Once the firm acknowledges its insolvency situation, its value may differ from the actual valuation of its assets, due to deadweight loss of different methods of resolution or to reputational effects even in the event of a bailout that prevents a liquidation. To make this explicit, we denote the value of the firm in states of insolvency as $\mathcal{V}_i(V(X, \theta))$ for any node $i \neq S$, with $\mathcal{V}_i(V(X, \theta)) \leq V(X, \theta)$. The dependence on i captures the fact that the ex-post value of the firm after resolution, $\mathcal{V}_i$, may be different depending on the method to resolve it. Payoffs in each node may depend as well on the method of resolution of the failing firm, as described below.

In the three nodes that correspond to resolution through OLA or liquidation or reorganization in bankruptcy, we consider the possibility that there may be negative externalities to society from the firm's failure. There is considerable debate in the academic and policy world about the existence, and importance, of these externalities. They are usually thought to originate from "spillover effects" or from "firesales." The label spillovers refers typically to credit disruptions: The inability of the failing firm to meet its debt obligations or perform functions that are essential to financial markets, such as being a broker-dealer (a "utility"), may negatively affect the ability to do business as usual for the firm's counterparties or clients. These externalities from spillovers, hence, are more likely to be important if the firm is a financial utility. The label firesales refers to the hurried sale of the assets of one large firm in distress at a price lower than its "fundamental" value—the price the firm could obtain if it had the time to find the right buyer. The externality would come from this sale unexpectedly lowering the market price of assets of certain types (those for which the failing firm had a significant market share). This could have the effect of decreasing the value of *other* institutions' balance sheets if they contain significant quantities of these undervalued assets. These externalities from firesales, hence, are more likely to be important in the market if the firm in distress is very large, or has a high market share for a particular asset.

Without taking a stand on whether these externalities are important, we include them in the model in order to study the implications that—were they shown to be significant—they would have in the process of financial firm resolution. We denote these potential externalities as $\alpha^i(X)$, for $i \in \{K_L, K_R, OLA\}$, with explicit dependence on the characteristics of the firm. Hence, we assume no externalities when the firm is solvent, receives a bailout, or is sold to a third party.

Next, we describe payoffs to each party in each of the six final nodes, as summarized in Figure 1. For each node, we also discuss the effect that having a regulated LW in place may have on the total value of the firm and on payoffs in that node.

Solvency $(i = S)$

First, we describe payoffs in the set of states when the firm is solvent, $\Theta^s(X)$, which correspond to the node $i = S$. In these states the value of the firm, net of taxes, $T(X)$, is larger than the value of debt, and thus creditors get repaid in full: $u_S^D(X, \theta) = D(1 + R)$. Welfare (the payoff to the planner) is $u_S^P(X, \theta) = u_i^F + u_i^D$, which equals the gross value of the firm, $V(X, \theta)$.

Crafting a good LW that conforms with regulatory requirements takes resources from the firm and limits its activities, decreasing its value. However, the control over more characteristics of the firm $\omega \subset X$ due to the LW review process may imply lower price of debt, which increases value. Hence, the overall effect of LWs on V in equilibrium is ambiguous.

Bailout $(i = B)$

If the firm fails, it may receive a bailout. The bailout is a transfer in the amount b from taxpayers to creditors and shareholders that allows the firm to repay its creditors in full and stay solvent. We denote the value of the firm after a bailout as $\mathcal{V}_B\left(V\left(X,\theta\right),b\right),$ to indicate that the transfer b will affect the final valuation of the firm; to keep the company solvent, the amount b will be such that

$$\mathcal{V}_B\left(V\left(X,\theta\right),b\right) - T\left(X\right) = D\left(1+R\right).$$

With this generic functional form, we recognize that the bailout allows the firm to continue its business, but the financial troubles revealed by the intervention could decrease the value of the firm. That is, we may have $\mathcal{V}_B\left(V\left(X,\theta\right),b\right)$ be lower than $V\left(X,\theta\right)+b$. The planner payoff is $u_i^P = u_i^F + u_i^D - \varepsilon_B b$, which equals $\mathcal{V}_B\left(V\left(X,\theta\right),b\right) - \varepsilon_B b$.

Because the firm is not liquidated in this node, having a regulated LW in place does not directly affect $\mathcal{V}_B$. The ambiguous effect on V, however, remains as described in the previous node.

Market Resolution $(i = M)$

A firm in distress may also be sold to the best buyer in the market for a value of $\mathcal{V}_M\left(V\left(X,\theta\right)\right).$ A scenario in which $\mathcal{V}_M\left(V\left(X,\theta\right)\right)$ is greater than the value of other methods of resolution corresponds to a situation in which the buyer may have a high valuation for the firm's assets, perhaps due to complementarities with the buyer's own assets. In this node, creditors get repaid in full. The value to society is $u_i^P = u_i^F + u_i^D$, which equals $\mathcal{V}_M\left(V\left(X,\theta\right)\right).$

A regulated LW can serve as a road map for interested buyers to calculate the full potential of the firm. Moreover, LWs imply that the set of choices x is more limited, which translates into a firm that is simpler to evaluate. Hence, having regulated LWs may increase $\mathcal{V}_M$ for a given $V\left(X,\theta\right).$

Bankruptcy for Reorganization $(i = K_R)$

If the firm files for bankruptcy, shareholders receive a 0 payoff, and debtholders get the value of the firm after reorganization, $\mathcal{V}_{K_R}(V(X,\theta))$ (they become the new shareholders).[5] The payoff to the planner takes into account potential externalities: $u^P_{K_R} = u^F_{K_R} + u^D_{K_R} - \alpha^{K_R}$, which equals $\mathcal{V}_{K_R}(V(X,\theta)) - \alpha^{K_R}$. The aggregate shock, θ^a, will affect the cost of DIP financing, $\bar{R}_2$, and hence the profitability of reorganizing the firm (Chapter 11) versus simply liquidating it (Chapter 7).

A regulated LW is crafted to include information that enhances the value of the firm in a bankruptcy procedure. It can also be valuable to external providers of DIP financing, since it makes the operations of the firm clearer, and supervision on ω gives a guarantee that the firm has desirable characteristics. Hence, a regulated LW increases $\mathcal{V}_{K_R}$.

Bankruptcy for Liquidation $(i = K_L)$

Shareholders again receive a 0 payoff, and debtholders get the liquidation value of the firm, $\mathcal{V}_{K_L}(V(X,\theta))$. The payoff to the planner takes into account potential externalities: $u^P_{K_L} = u^F_{K_L} + u^D_{K_L} - \alpha^{K_L}$, which equals $\mathcal{V}_{K_L}(V(X,\theta)) - \alpha^{K_L}$. Choosing liquidation over reorganization because of the absence of cheap DIP financing could imply some inefficiencies, since in reorganization the firm could continue to operate for a time after failure. This could avoid the inefficient liquidation of specific assets and the termination of otherwise valuable derivatives contracts. However, whether $\mathcal{V}_{K_R} > \mathcal{V}_{K_L}$ will also depend on the realization of θ^a through the price for DIP financing. In the same spirit, assuming that reorganization would call for fewer contracts terminated suddenly, and hence less disruption in the market, we assume externalities are less important under reorganization: $\alpha^{K_L} > \alpha^{K_R}$.

For the same reasons listed for the node of bankruptcy for reorganization, a regulated LW increases $\mathcal{V}_{K_L}$. Possibly, however, the positive effect in $\mathcal{V}_{K_R}$ for a given $V(X,\theta)$ is larger than the positive effect in $\mathcal{V}_{K_L}$: With an LW the price of DIP financing may be lower, but under bankruptcy for liquidation the firm does not look for financing, so this relative advantage of the LW is irrelevant.

[5] It is possible that after a firm files for bankruptcy (or even under other methods of resolution) shareholders get positive value after liquidating the firm and repaying debtholders. However, this depends on further resolution of uncertainty about the actual price that the market pays for the assets of the liquidating firm, which we are not modelling here. For the purpose of our analysis, the relevant value is $\mathcal{V}$, which can be interpreted as the expected value after resolving the firm, calculated at the moment when the planner or shareholders need to choose the method of resolution. By the definition of insolvency, $\mathcal{V} \leq V$.

OLA $(i = OLA)$

If OLA is used instead of bankruptcy to resolve the firm, shareholders again receive a 0 payoff, and debtholders get the liquidation value of the firm, denoted $\mathcal{V}_{OLA}(V(X,\theta), b)$. The payoff to the planner takes into account potential externalities, as well as the cost of injecting funds b to finance the wind down of the firm: $u^P_{OLA} = u^F_{OLA} + u^D_{OLA} - \alpha^{OLA} - \varepsilon_{OLA} b$, which equals $\mathcal{V}_{OLA}(V(X,\theta), b) - \alpha^{OLA} - \varepsilon_{OLA} b$. The value of a firm resolved through OLA may be different than $\mathcal{V}_{K_L}(V(X,\theta))$ or $\mathcal{V}_{K_R}(V(X,\theta))$. We assume $\mathcal{V}_{OLA} > \mathcal{V}_{K_L}$ in bad aggregate states of the economy to capture the possibility that cheap interim financing is available from the OLF, rather than having to depend on the availability of DIP financing in the market. We also assume that $\alpha^{OLA} < \alpha^{K_R}$ (less externalities under OLA than under bankruptcy) to reflect the fact that institutional advantages of OLA such as the ability to impose two business days of automatic stay even to qualified financial contracts (QFCs) may prevent some of the firesale effects on other firms.[6]

A regulated LW will help increase $\mathcal{V}_{OLA}$ in the same way as in bankruptcy, providing useful information. However, the increase in value will be more limited than for reorganization, since the interim financing needed to efficiently wind down the firm is not DIP financing (the price of which would be sensitive to that information), but rather funds provided by the OLF.

The payoffs to the players in each of the nodes can be summarized as follows:

- The payoff to the shareholders of the firm, u_F, will be the value of the firm, contingent on solvency, net of repayment of the debt and tax obligations:

$$u_F = \begin{cases} V(X,\theta) - D(1+R) - T(X) & \text{for } i = S. \\ 0 & \text{for } i \neq S; \end{cases}$$

- The payoff to creditors, u_D, will be the principal of the debt, D, plus interest, whenever the firm is solvent; if the firm is insolvent, they will get the resolution value of the firm, which will depend on the alternative chosen to resolve it:

$$u_D = \begin{cases} D(1+R) & \text{for } i = S \\ \min\{D(1+R), \mathcal{V}_i(V(X,\theta), b)\} & \text{for } i \neq S. \end{cases}$$

[6] In December 2014, in recognition of the inefficiency of the early termination of swap contracts at the moment when insolvency is determined, 18 major global banks agreed to a protocol developed by the International Swaps and Derivatives Association to include temporary stays on termination rights. The objective of the agreement was to give bankruptcy judges or regulators in charge of resolution time to conduct an orderly liquidation, since most commentators agreed that even two extra business days under OLA would not be enough to prevent inefficient liquidations. See Gruenberg (2015).

- The payoff to the social planner, u_P, will consist of the value of the firm, which will vary depending on whether it is solvent or, if insolvent, the method of resolution. If the firm is resolved through bankruptcy or OLA, externalities may decrease the payoff. If there is a transfer of taxpayer funds or from the OLF, there will be a cost determined by a realization according to $g(\varepsilon)$:

$$u_P = \begin{cases} V(X,\theta) & \text{for } i = S \\ \mathcal{V}_i\left(V(X,\theta), b\right) - \alpha^i(X) - \varepsilon_i b, & \text{for } i \neq S, \end{cases}$$

where

$$\alpha^i(X) \begin{cases} > 0 \text{ for } i \in \{K_L, K_R, OLA\} \\ = 0 \text{ otherwise,} \end{cases}$$

and

$$\varepsilon_i \begin{cases} > 0 \text{ for } i \in \{B, OLA\} \\ = 0 \text{ otherwise.} \end{cases}$$

It is worth pointing to two assumptions we are making in the set up of the game. First, we assume that the firm cannot misrepresent its choices in X, i.e., regulators get in the LW a truthful and accurate description of the firm and its potential resolution strategy. Second, we are assuming that all players know the value of all the final payoffs of the game, including the potential externalities and the different values that can be achieved depending on the method of resolution of a firm in financial distress. These are not realistic assumptions. We make them here to present the simplest environment that allows us to highlight the main economic forces behind the commitment problem in financial firm resolution. Relaxing them will certainly complicate the analysis of the problem, but we believe that the spirit of our conclusions in this article will remain valid in a more realistic environment with asymmetric information and uncertainty about payoffs.

It is our objective to understand the value that LWs bring to the problem of SIFI regulation, in particular how they may help to provide commitment to the policymaker not to choose a bailout in the event of insolvency. To that end, in the next subsection we characterize the outcome of the game, and in Section 2 we discuss how this outcome will differ with and without regulated LWs.

Outcomes

Given the sequential moves in the game, we look for a subgame perfect Nash equilibrium. The set of equilibrium strategies $(\sigma_F^*, \sigma_D^*, \sigma_P^*)$

are such that, at any decision node, agents are maximizing their expected value of the corresponding subgame. The first relevant subgame, working backward from the end nodes, is that of the firm choosing $a_F \in \{M, K_L, K_R\}$ if the planner has chosen not to intervene; that is, the firm chooses between selling to the market, liquidation, or reorganization. In this case, for each given vector of possible firm characteristics, (ω, x, R, a_P), and uncertainty realization up to that node, (θ, ε), the firm simply chooses the strategy that maximizes the value of the firm:

$$\max_{a_F \in \{M, K_L, K_R\}} \mathcal{V}_{a_F}(V(X, \theta)).$$

This subgame is then substituted by the final payoffs corresponding to this optimal choice of the firm, denoted $a_F^*(\omega, x, R, a_P; \theta, \varepsilon)$. This creates a new final subgame in which the planner chooses $a_P \in \{NI, B, OLA\}$; that is, it chooses from no intervention, an intervention in the form of a bailout, or resolution through OLA. To make this choice, it compares the value under each feasible method of resolution, for each realization of θ, taking a_F^* as given:

$$\max \left\{ \begin{array}{c} \mathcal{V}_{a_F^*}(V(X, \theta)) - \alpha^{a_F^*}(X), \mathcal{V}_B(V(X, \theta), b) - \varepsilon_B b, \\ \mathcal{V}_{OLA}(V(X, \theta), b) - \alpha^{OLA}(X) - \varepsilon_{OLA} b \end{array} \right\}.$$

This optimal choice is denoted $a_P^*(\omega, x, R; \theta, \varepsilon)$. In the next relevant subgame, creditors choose R anticipating the subsequent choices of intervention and resolution (that is, a_F^* and a_P^*), and using the probability of failure contingent on a choice of the firm, $\pi(X)$, to calculate their expected payoffs. That is, they set $R^*(X)$ given a_F^*, a_P^*, to solve

$$E_{\theta,\varepsilon}[u_D(R, a_P^*, a_F^*); X] = D(1 + \bar{R}_1), \quad (1)$$

where $E_{\theta,\varepsilon}$ denotes the expectation taken with respect to the densities $h(\theta)$ and $g(\varepsilon)$. In the next relevant subgame, the firm chooses its strategy $x^*(\omega)$ anticipating the interest R^* and the choices of intervention and resolution a_F^* and a_P^*, and again using $\pi(X)$ to calculate their expected payoffs:

$$\max_x E_{\theta,\varepsilon}[u_F(x, R^*, a_P^*, a_F^*; \omega)];$$

that is, the firm chooses x, the elements of the firm characteristic vector that are not subject to the constraints imposed by supervisors according to safety and soundness measures, or the LW review process. Finally, in the last subgame, at the initial node, the planner chooses regulation ω^* (the limits on the regulated characteristics of the firm) anticipating the choices at subsequent nodes, x^*, R^*, a_P^*, and a_F^*, as well as the corresponding $\pi(X^*)$:

$$\max_\omega E_{\theta,\varepsilon}[u_P(\omega, x^*, R^*, a_P^*, a_F^*)].$$

2. THE TIME INCONSISTENCY PROBLEM

We are now ready to formally describe the time inconsistency problem behind the TBTF problem. If commitment was available, the planner would like to choose, at time zero, both regulation, ω^*, and resolution policy, a_P^*, to maximize expected payoff

$$\max_{\omega, a_P} E_{\theta,\varepsilon}\left[u_P\left(\omega, x^*, R^*, a_P, a_F^*\right)\right].$$

Note that in the equilibrium we just described, since commitment is not available, at time zero the planner chooses only ω, taking as given his future subgame perfect choice a_P^*. Most importantly, x and R are chosen by the firm and the creditors also using a_P^*.

It is easy to see from the game tree that the solutions with and without commitment will not coincide. From the payoffs of the debtholders, we see that they only take losses if the firm is insolvent and resolved through bankruptcy or OLA. That is, bailouts guarantee debtholders get repaid in full. Because of this, a firm will find that it can minimize the interest rate on its debt, R^*, if it chooses a set of characteristics in x that make it profitable when solvent but likely to fail (i.e., "risky"), provided these characteristics also make it hard to resolve through bankruptcy or OLA. This is the moral hazard problem (i.e., choosing an x that implies too much risk of insolvency given the costs to society of a failure) that is triggered by TBTF (i.e., the firm being too hard to resolve). Excessive risk-taking of the firm is optimal because debtholders will not demand a higher interest rate, R, to compensate for it, since their payment is likely to be guaranteed by a bailout.

Instead, implementing a policy that prohibits bailouts (i.e., committing ex ante to never choosing $a_P = B$ for any X and θ) lowers the probability of financial distress because debtholders stand to lose from firm failure and hence demand a high R for x choices that make the firm likely to fail (see equation 1). This sensitivity of R to x choices implies that the firm finds it profitable to choose an x that makes it less likely to fail.

The time inconsistency problem arises because it is not credible for the planner to commit to never choosing $a_p = B$, since this is the optimal ex-post choice for some combinations of X and θ—for example, when the aggregate state of the economy is bad and DIP financing is so expensive that it makes the firm fail or when externalities are high. Because of this, the planner would need some external commitment mechanism in order to not choose to bailout in the relevant subgame. This issue of commitment arises because bailout policy is only implemented at the time of financial distress (that is, after observing the realization of θ). Since at that point (in that subgame) the characteristics of the firm, X, have already been determined, the planner finds it profitable to bail out the firm whenever the payoff to society is higher with intervention than without it. This implies $a_P = B$ in a larger set of states θ than it would result under commitment. That is, denoting with b_B and b_{OLA} the amount of funds needed to bail out the firm or resolve it through OLA, respectively, a bailout will be chosen by the policymaker whenever

$$\mathcal{V}_B\left(V\left(X,\theta\right),b_B\right)-\varepsilon_B b_B > \max\left\{\begin{array}{c}\left\{\mathcal{V}_i-\alpha^i\left(X\right)\right\}_{i\in\{K_L,K_R,M\}}, \\ \mathcal{V}_{OLA}-\alpha^{OLA}\left(X\right)-\varepsilon_{OLA}b_{OLA}\end{array}\right\}. \quad (2)$$

This behavior is anticipated by the firm in equilibrium, making the announcement of no bailouts at time 0 irrelevant. That is, $a_P\left(X,\theta\right) \neq B$ for all X,θ is a time-inconsistent strategy. This inability of the planner to commit will make the firm more likely to choose certain x characteristics that will imply more frequent failure. The reason is that for some of these failures creditors will be bailed out, and hence R for those x choices will be lower than it would be under commitment, according to the equilibrium interest rate condition in equation 1. Unless externalities are very large, the equilibrium under commitment would result in higher expected welfare by eliminating moral hazard in the choices of x. In the next section, we discuss specific examples of the manifestation of this moral hazard problem in the choices of the firm.

How can LWs help provide commitment not to bail out? Because the existence of a regulated LW implies higher values for the firm on the right-hand side of equation 2 but not on the left-hand side, the set of states for which the inequality will be satisfied will be smaller when the LW is in place. That is, by controlling certain firm characteristics that are mostly relevant for ease of resolvability, such as size and complexity, and by spelling out strategies to maximize the value of the firm in the event of liquidation, LWs make the alternative of unassisted resolution more attractive to the planner when compared to a bailout in the ex-post event of firm insolvency.

3. FIRM CHOICES AND MORAL HAZARD

To understand the implications of the design of regulation and resolution methods on outcomes, it is useful to be more specific in describing firm choices. In what follows, we provide a (nonexhaustive) list of salient firm choices that we had summarized in X in the description of the game above. We also discuss how each may influence the payoffs and probabilities of failure in the game. We consider choices over the following characteristics: the firm's size, its production and legal structure, its risk choices, the liquidity of its asset portfolio, its leverage, and the degree of maturity transformation and other characteristics of its financing through debt. We discuss each of these in detail next.

Size

We denote the size of the firm by A. The firm benefits from increasing its size because of economies of scale, that is, $V(X, \theta)$ increases with A. The larger the firm, however, the more difficult it is to have an orderly liquidation in the event of financial distress. That is, the loss in value from resolution, $V - \mathcal{V}_i$, is increasing in A for $i = M, K_R, K_L, OLA$. It is plausible that this loss is greatest for K_L, which implies the fastest liquidation. Similarly, the externality cost, $\alpha^i(X)$, is increasing in A for $i = K_R, K_L, OLA$, and we expect it to be the largest for K_L and the lowest for OLA, since this institution can impose a two-business-day stay on qualified financial contracts, preventing some firesales of assets. This means that the failure of a larger firm may be more likely to trigger intervention (i.e., the planner prefers B to OLA, and OLA to NI), since it has the potential to impose larger externalities on the economy if it fails. In summary, an increase in size lowers the expected repayment of creditors in the nodes K_L, K_R, M, and, to a lesser extent, in OLA, but it increases it in the nodes of S, B; moreover, the node of bailout, B, in which creditors get repaid, happens with higher probability.

Complexity in Production

We denote the level of complexity due to production complementarities as C^p. This captures the number and interconnectedness of subsidiaries. The firm benefits from increasing C^p because of economies of scope if it stays solvent, i.e., V is increasing in C^P. However, if the firm fails, more complexity will make assessments of the value of the assets difficult unless the firm is bailed out and it continues to operate as usual, i.e., $V - \mathcal{V}_i$ is increasing in C^P for $i = M, K_R, K_L, OLA$.

For a firm of high complexity, and all else equal, OLA and bankruptcy may be at a disadvantage with respect to the market (other financial institutions may be more experienced at doing due diligence, and they may even be counterparties of the failing firm, who are better able to evaluate its portfolio). The OLA may be better than a bankruptcy judge, since the regulators have a good deal of information about the firm that may help them liquidate the firm, preserving some of the economies of scope. The value of bankruptcy through liquidation may also be higher than in reorganization, since complexity will again interfere with due diligence necessary to obtain DIP financing, making it expensive. Hence, a plausible ranking of payoffs for a firm of very high C^P would be

$$\mathcal{V}_B > \mathcal{V}_M > \mathcal{V}_{OLA} > \mathcal{V}_{K_L} > \mathcal{V}_{K_R}.$$

Complexity Due to Regulatory Arbitrage

We denote the level of complexity due to regulatory arbitrage as C^b. It represents the complexity in firm structure due to having subsidiaries in different countries or even simply having legal entities that do not coincide with business units. Even though it may be convenient for the firm to exploit different tax regimes across different borders, this implies it is subject to different regulatory environments and legal systems that would handle bankruptcy in potentially different manners. This implies that maximizing the value of the firm in resolution would require a level of coordination across jurisdictions that seems difficult to attain. Hence, it is likely that in states of failure without a bailout, the effects of this type of complexity on payoffs are similar to the ones we just described for C^P, i.e., we have that $V - \mathcal{V}_i$ is increasing in C^b for $i \neq B, S$. However, the effect on welfare from an increase in C^b when the firm is not resolved is different than that of an increase in C^P. In contrast with savings that arise with economies of scale or scope, tax savings for the firm are simply a transfer between the rest of the economy and the stakeholders of the firm, so the planner does not value them. Moreover, they may actually distort the decisions of the firm, making its value—before taxes—actually lower for society. This is easily seen by comparing the payoff to shareholders in states of solvency, $V(X, \theta) - D(1 + R) - T(X)$, with the payoff of the planner, $V(X, \theta)$. Clearly, because of the dependence of the tax bill on X through C^b, the level of complexity that maximizes the expectation over the value of the firm net of taxes may not coincide with the one that maximizes the expectation over the gross value of the firm, which is what the planner cares about.

Balance Sheet Risk

To explain the risk-shifting problem that arises in the financial system due to explicit and implicit government guarantees, we consider the firm's choice of risk in its balance sheet. For this, we assume that financial assets can be classified into "risky" and "safe," and the firm choices in X include the proportion of risky assets in the balance sheet of the firm, denoted as $\rho \in [0,1]$. Assuming that the returns of both assets are stochastic and equal to r_R and r_S, correspondingly, they imply a return on the balance sheet that depends on ρ:

$$\rho r_R + (1-\rho) r_S.$$

A risk-shifting problem arises when the risky project is not efficient to undertake for society or creditors. A sufficient condition for that situation would be $E[r_R] < E[r_s]$. Under this assumption, despite it being inefficient, limited liability may still make it attractive for shareholders: Given that shareholders are not liable for losses in the event that the realized return is lower than the cost of the assets, we have that for them

$$E[\max(r_R, 0)] > E[\max(r_S, 0)].$$

Lenders, on the other hand, suffer losses but do not get an upside when risky projects pay off. Hence, after they observe the choice of ρ of the firm, they adjust the price of funding for the assets, R, to break even given the choice of risk of the firm.[7] Through this mechanism, debt monitoring implies that risk-taking is costly for the firm, and hence it does not choose only risky projects. Importantly, though, the payoffs to creditors are guaranteed to be $D(1+R)$ whenever there is a bailout. This implies that the sensitivity of the price of debt, R, to the risk choices of the firm is decreased if lenders believe that financial distress is likely to lead to the planner's intervention through a bailout. In turn, this means a higher risk chosen by the firm, which translates into higher probability of financial distress, $\pi(X)$.

Note that, short of the complete guarantee of debt that follows from a bailout in our simple model, anything that increases the value of an insolvent firm will increase the expected payoff to creditors. For example, if resolution through OLA is likely to enhance the liquidation value of the firm because of institutional advantages such as the ability to impose a two-day stay on qualified financial assets, or because of the availability of interim financing provided by the OLF, this will have the effect of decreasing the sensitivity of R to risk choices.

Liquidity

We denote the fraction of the portfolio that is invested in liquid assets as $\lambda \in [0,1]$. A high λ has several (and potentially opposing)

[7] In reality, lenders are likely to observe risk choices of the firm they lend to only imperfectly. As long as they have a (possibly costly) way of gathering some signal about this choice, the intuition conveyed in our simple model would still hold.

effects. First, it decreases the probability that the firm gets in financial distress because of liquidity needs: If, due to market frictions, the financing ability of the firm deteriorates (for example, it is unable to roll over short-term debt or unable to access repo markets), the firm can easily deplete its stock of liquid assets, remaining solvent along the way. However, this comes at a cost: Liquid assets, such as Treasury securities, have a liquidity premium, which means their return is lower. This means that a second effect is lower margins, which makes the firm less resilient to shocks and more likely to fail. High liquidity, however, has value in the event that resolution becomes necessary, since assets that are liquid are typically valued by a large number of agents in the economy, and their value is independent from the rest of the portfolio of the firm that holds them. Hence, a third effect is that a high λ is likely to imply a high liquidation value irrespective of the method of resolution. A fourth effect follows: A high resolution value implies that the creditors of the firm expect to recover a larger amount in the event that the firm goes into bankruptcy, and hence the price of debt decreases. If these savings are large enough, they may compensate the liquidity premium paid to achieve a higher λ. Finally, from the perspective of the policymaker, a fifth effect may arise: If the firm holds a large quantity of liquid assets, which are likely to be also in the balance sheets of other financial institutions, this may trigger a concern about firesale effects in the event of liquidation. Hence, indirectly, a higher liquidity requirement may make methods of resolution that are believed to prevent firesales, such as a bailout or OLA, more attractive.

Leverage

We denote the proportion of assets that are financed with equity capital instead of debt as κ. Leverage is $\frac{D}{V} = 1 - \kappa$. If κ is larger, the asset value realizations for which the firm is insolvent decreases. Capital requirements of the form $\kappa \geq \underline{\kappa}$, hence, decrease the probability of failure π. Tax advantages will, for the same cost of capital and debt, make the firm prefer leveraging. Only a high price of debt will make the firm choose $\kappa > \underline{\kappa}$ in the absence of regulatory constraints. However, a firm with a high κ should face a lower probability of failure, which will in turn lower its cost of debt, increasing V.

In the event that the firm does go into failure, higher κ will imply less disruptions for the economy in the form of contagion through counterparty risk, making $\alpha^{K_L} - \alpha^{OLA}$ smaller. Also, less reliance on debt will probably mean less fragile debt, implying less need for DIP financing. This will make the difference $\mathcal{V}^{OLA} - \mathcal{V}^{K_L}$ smaller. Hence, higher capital requirements will undermine OLA's advantage with respect to bankruptcy.

Fragile Debt and Maturity Transformation

We denote as δ a measure of the reliance of the firm on "fragile" debt and the degree of maturity transformation that it engages in.[8] Fragile debt includes both run-prone debt financing (short-term debt like commercial paper and unsecured deposits) and repos and other QFCs.

Standard priority rules in bankruptcy—given by deposit insurance, bankruptcy law, exemption from the automatic stay for QFCs, as well as private arrangements between creditors and SIFIs (i.e., senior and junior debt denominations)—determine the probability that fragile-type debt gets repaid in the event of liquidation. For example, deposit insurance implies that depositors always get repaid. This implies that such depositors will only demand the risk-free rate of return, and hence the price of these loans will remain unchanged with changes in the firm's choices of X. Less explicit arrangements, such as the safety net implied by the existence of a policymaker with funds available to finance a firm in trouble, imply that short-term debt that is believed to be a mechanism for the contagion of financial weakness will be more likely to be first in line for repayment in a liquidation. Moreover, short maturity means that when information first starts to appear about financial trouble for a firm, its short-term lenders are simply able to not roll over their loans. These reasons imply that the lenders who own the fragile debt expect to get repaid with high probability, and hence this type of debt has a lower cost for the firm.

Relying on short-term debt has two important negative effects on the strength of the firm. First, for a given capital structure, more fragile debt (higher δ) implies a higher probability of financial distress. Because of its short maturity, this debt needs to be refinanced frequently and hence is more likely to become unavailable when financial weakness appears, making insolvency more likely. To limit this problem, safety and soundness regulations impose certain requirements, such as the liquidity coverage ratio (LCR) recently established in the international Basel III accord, which requires a certain balance between the liquidity of the assets of a firm and its fragile debt. Second, higher δ implies higher effective DIP financing needs of the firm to continue business

[8] The financial institution engages in maturity transformation when it borrows short-term (e.g., through deposits and repurchase agreements) and lends long-term (e.g., through mortgages and industrial loans), acting as an intermediary between lenders and borrowers in the economy.

as usual in the event of bankruptcy. This is mainly because once the firm has failed, any QFCs are exempt of the automatic stay, and their liquidation means that counterparties retain collateral, depleting the assets of the firm at a discount rate (determined by the haircut on the collateral).

The purpose of targeting a high LCR is to tame the first negative effect, i.e., decrease the probability of failure. The purpose of using the LWs process to quantify and minimize $\delta(1-\kappa)A$, instead, is to minimize the second negative effect, i.e., limit the DIP financing needs. In summary, the traditional approach of safety and soundness regulation differs from, and is complemented by, the LWs review process.

This concludes our description of the choices of the firm. Summarizing, in our model, the vector X consists of these key firm characteristics:

$$X = \{A, C^p, C^b, \rho, \lambda, \kappa, \delta\},$$

namely, the size of its balance sheet, complexity in production, complexity due to regulatory arbitrage, balance sheet risk, liquidity, equity financing, and the structure of its debt. It is plausible to think that, while characteristics such as equity financing and liquidity choices have been highly scrutinized by supervisory controls since before DFA, they are now being evaluated under a new perspective: what choices are likely to facilitate the orderly liquidation of the firm in case of distress—as opposed to what choices are going to minimize the probability of distress. Other characteristics, such as size of balance sheet, complexity in production, complexity due to regulatory arbitrage, or the structure of debt, although always important to supervisors, are now formally required within the LW review process to comply with the standard of facilitating resolution. This constitutes an expansion of the subset ω of firm characteristics that are constrained by regulation because of the introduction of LW requirements under Title I of DFA. Finally, characteristics that are harder to measure and control, such as the risk-taking of the firm, are at the heart of the choices of the firm x that remain subject to the moral hazard problem with TBTF institutions.

How Living Wills May Help With Time Inconsistency

Reflecting regulatory requirements in DFA, we have modeled LWs as instruments to *disclose* the efficient way to liquidate a firm, as well as to *control* a subset of firm characteristics ω in X. In reality, the way in which regulators control firm characteristics is through an iterative process with the regulated firms. That is, if a firm's choices in ω described in the LW imply that its unassisted resolution plan is deemed "noncredible" by the regulators, the firm is required to change its choices in ω in order to solve the shortcomings of its resolution plan. For example, if the LW describes a firm that is too large or too complex, it may be asked to divest some of its assets. Or, if the proportion of fragile debt is too large, regulators can ask the company to change their capital structure to compensate for this added fragility. Then, a new LW is crafted given the new structure of the firm, and it is evaluated again by regulators.

As we argued in Section 2, regulators hope that LWs will alleviate the time-inconsistency problem of the planner. They may do this by minimizing both the intensive (the amount b) and extensive (the probability π) margins of intervention. They accomplish this in two related ways:

1. Ex post: by disclosing information on how to maximize the liquidation value of the assets of the firm;

2. Ex ante: by controlling the choices of the firm in ω and disclosing that an efficient resolution plan is in place that does not involve intervention.

As we argued when discussing payoffs in Section 1, the ex-post gains occur because a good LW maximizes the liquidation value of the assets, since (i) it constitutes a detailed description of the assets and business model of the firm (e.g., by describing economies of scope, location, and logistic needs of core functions); (ii) it minimizes the market disruptions triggered by the failure of the firm (e.g., by listing the main counterparties and any relevant cross guarantees); and (iii) it increases the likelihood of market-based reorganization options (e.g., by providing a readily available description of financing needs and better information about the company that implies better pricing of DIP financing, higher likelihood of a private acquisition, or less need of assistance in a purchase and acquisition process). This implies that LWs can improve the payoff to society, u_P, in the nodes that do not involve the planner's intervention, $i = M, K_R, K_L$, by increasing the liquidation value of the firm, $\mathcal{V}_i\left(V\left(X, \theta\right)\right)$, and/or reducing the externalities $\alpha^i\left(X\right)$.

The requirement of LWs will bring ex-ante gains to society whenever the increase in payoffs, u_P, in the no-intervention nodes implies that the strategy of the planner of no intervention dominates that of intervention, as indicated in equation 2. When the firm anticipates fewer instances in which a bailout will be chosen given insolvency, it finds it profitable to make choices in x that decrease the probability of insolvency, $\pi\left(X\right)$. This way, LWs may lessen the TBTF problem.

Related Proposals

There are several regulatory proposals and initiatives that relate to the LW's power to commit not to bail out. For the purpose of our discussion, they can be classified in two groups.

The first group of proposals seeks to enhance clarity and commitment to priority rules that assign losses in bankruptcy. For example, Barry Adler of NYU Law has presented the idea that LWs can effectively establish a priority rule in case of failure; that is, it would indicate how debt of different classes (or "priority tranches") would be converted sequentially into capital if the firm is in financial trouble.[9] This way the firm would automatically recapitalize until it is solvent again. A clear benefit is that a LW would make transparent the incentives of the owners of more junior debt to make the interest rate they demand depend on the ex-ante risk choices of the firm. Moreover, with iterative conversion of debt into equity preventing failure, it could minimize the liquidation of collateral, since the most senior debt may still be repayable.

In a similar spirit are proposals to impose certain requirements on the combination of capital and long-term debt ("total loss capacity") at the parent of a bank holding company (BHC). The idea is that these debt holdings at the parent company would serve as a cushion for losses in subsidiaries, which are in turn financed by the parent company holding their equity. This structure would be particularly useful if the BHC were to be resolved using a strategy of "single point of entry" (where only the parent company files for bankruptcy), as favored by the chairman of the FDIC in recent speeches (see Gruenberg 2015). Herring and Calomiris (2011) had previously argued for a similar cushion for losses in the form of convertible debt. This is debt that converts automatically to capital under certain triggers tied to financial strength measures. These proposals, as well as any other effort that selectively puts creditors who are able to monitor the firm's choices in X on the hook for losses in the event of financial distress, could complement the requirement and regulation of LWs in the objective of alleviating the TBTF problem.

A second group of proposals seeks to enhance the efficiency of unassisted bankruptcy. For example, an important effort has been made to make financing through "swaps" less problematic in the event of failure. A swap is a type of derivative, a financial contract between two parties that agree to exchange cash flows replicating the payments of underlying securities (for example, the coupons of a bond or a future

[9] See Chapter 8 in Acharya et al. (2010).

on an exchange rate). Changes in the value of these swaps are typically backed by posting collateral (i.e., if a change in the market implies that counterparty A is expected to have to transfer to counterparty B a cash flow with net present value of M, then to insure B against the failure of A we will see A transferring collateral with value M to B). In the event of bankruptcy filing, these swap contracts are qualified to be exempt of the automatic stay. This means, in our example, that if A files for bankruptcy, the swap is automatically terminated, and B keeps the collateral amount M—hence depleting B of these funds immediately. In a recent joint initiative of the FDIC and the International Swaps and Derivatives Association, 18 of the largest global bank holding companies, which constitute a majority of the swap market, have voluntarily agreed to end the automatic termination of covered derivative contracts in the event of a bankruptcy or public resolution of a systemic financial institution, effectively changing the priority ordering of creditors. This should help minimize the need of DIP financing of a bankrupt firm, increasing its liquidation value. Similar agreements, if they were to include other types of derivatives, such as repurchase agreements, would erode one of the main advantages of OLA with respect to bankruptcy: the brief two-business-day stay that OLA can impose on all qualified financial instruments.

Researchers at the Hoover Institution Resolution Project have proposed more encompassing reforms of bankruptcy law that would go beyond curtailing exemptions to the automatic stay, giving increased flexibility to judges and even allowing for external interim funding as part of the procedure (Jackson 2014). These proposals are compatible, and would work best, with a regulated LW that provides useful information to bankruptcy judges, as mandated by DFA.

4. IMPLICATIONS

We have so far answered our main question of how can LWs help tame the moral hazard problem behind TBTF. In this last section we focus instead on two related questions. First, we are interested in learning what are the contexts in which requiring LWs is more likely to be helpful. Second, we would like to know what are the key characteristics that make LWs the most useful. As a preview of our conclusions, we find that LWs are most likely to be useful in the absence of other forms of commitment not to bail out, if externalities are important, and whenever institutions are in place such as the OLA (with the ability to more efficiently finance wind downs, which provides higher value to unsecured lenders). We also find that LWs are most useful when they communicate enough information to creditors about their expected losses in the event of failure.

When is Regulation of LWs Likely to be Most Useful?

A good LW provides information that maximizes the going concern of the failing firm. Because of this feature, one could think about LWs as a regulated version of debt covenants or credit line provisions, which are put in place by market participants in the absence of a safety net or financial firm regulation. Because creditors increase their expected payouts if the going concern of the firm is maximized in the event of failure, structuring the firm in a way that makes it easy to resolve will increase the liquidation value of the firm, and hence the expected payoff to creditors in the event of failure.

In this section we will compare outcomes for society depending on whether LWs are unregulated (in the form of debt covenants) or regulated (supervisors review them). Why and when is the regulation of LWs needed? There are two important frictions why the market may not impose the same constraints on firms under privately arranged debt covenant than under regulated LWs: lack of commitment and the existence of externalities from liquidating a SIFI. We discuss them next.

Lack of commitment implies too many failures and makes the regulation of LWs necessary. Our model illustrates that, in the absence of commitment to not bail out, the price of debt may be insensitive to the quality of the market-provided LW (or, in other words, to the expected liquidation value of the firm in the event of insolvency). In other words, the price of debt may be hump-shaped in firm characteristics such as size, complexity, and risk, since a bailout is expected when these take high values. This results in market-provided LWs that do not impose the right limits on ω and firm choices that imply a high probability of a bailout (firms that are too large, complex, illiquid, and risky). The need to require credible LWs (i.e., limits on ω that are regulated) parallels (and complements) other safety and soundness regulations, such as capital requirements, in the presence of a safety net for financial firms. Summarizing:

Implication 1 *Whenever the planner cannot commit to a resolution strategy, unregulated LWs are not efficient, and the probability of default is inefficiently large.*

Externalities also make the regulation of LWs necessary. Our analysis illustrates that, even if the planner would have the technology to commit to never intervene, the quality of the LWs may be suboptimal if there are externalities. This is because, even though the price of debt does respond to the limits on X imposed by the unregulated LWs in this hypothetical scenario without bailouts, creditors do not internalize the spillovers to the rest of the economy of the firm failing, resulting in suboptimal limits on X for unregulated LWs. In fact, in this hypothetical scenario without bailouts, externalities are the only reason for regulating LWs. Summarizing:

Implication 2 *Whenever the planner can commit to a resolution strategy, unregulated LWs are efficient if and only if externalities are not important.*

Our analysis also points out that, even in the absence of the two frictions we just discussed, regulating LWs may not be enough to achieve efficiency. In particular, we can think of two instances in which welfare can be improved ex post by having an institution provide interim financing to wind down insolvent SIFIs. OLA is one such institution. First, if aggregate macroeconomic conditions are bad the cost of DIP financing will be high. OLA has an advantage in these instances, since it can use the OLF to fund the firm's operations. This will mean firms that fail in situations of economic crisis may be more efficiently resolved through OLA. Second, if externalities are important, the greater likelihood of an orderly liquidation through OLA than bankruptcy will also mean that the availability of OLF funds may improve ex-post welfare. If neither of these two conditions are present, however, LWs can provide the same welfare relying solely on bankruptcy. Summarizing:

Implication 3 *OLA may improve efficiency by providing financing below DIP market rates in states of bad aggregate conditions or whenever externalities are important.*

A different argument in favor of OLA that is commonly cited relies on the fact that it makes DIP financing available for a firm facing a "pure liquidity shock." For such a firm, liquidation would be inefficient, and the availability of OLF funds could improve welfare. For this argument to be relevant, however, one would need to justify why solvent firms may become illiquid.

We have identified conditions under which OLA may be needed on top of regulation of LWs to improve ex-post welfare. A caveat arises when analyzing the equilibrium of the game, however: OLA needs credible LWs to improve ex-ante welfare. This is because the same cheap DIP financing that improves welfare ex post by increasing the liquidation value of the firm will make failure less costly to creditors of such firms. Hence, commitment not to bail out will be sufficient for the existence of OLA to be welfare enhancing, but in its absence the positive effect ex post may be cancelled by an increasing probability of failure, π, through the moral hazard in the choices of x. Regulating the quality of LWs improves liquidation outcomes, but also, importantly, it expands the set of characteristics ω on which regulators put constraints, decreasing the incentives of the planner to intervene ex post. In other words,

Implication 4 *Whenever the planner cannot commit to a resolution strategy, for OLA to increase ex-ante welfare we need to complement it with regulated LWs that (i) improve liquidation outcomes outside of OLA ($i = M, K_L, K_R$), and (ii) constrain the choices of the firm.*

These results raise some important issues. First, how are regulators supposed to identify being in an aggregate state that makes funds for market-based resolution scarce? If such a state can be determined following objective criteria, should OLA only be invoked using such an objective trigger? Second, if the main contribution of OLA is to provide funding, should this institution be defined in this narrower way, leaving the actual reorganization of the firm to a bankruptcy court, which is more likely to respect priority rules?[10] Limiting the powers of regulators to intervene in the event of financial distress will decrease moral hazard, but our analysis indicates that setting the limits right is a complicated matter.

We have established that LWs can help to increase ex-ante welfare given the existence of OLA and the possibility of other forms of bailouts. It is important to emphasize as well that the optimal structure of the firm balances efficiency in normal times with low likelihood of failure *and* ease of resolution. More information on externalities and evidence on the ability of different methods of resolution to maximize the liquidation value of firms would be valuable input to improve the regulation of financial firm resolution.

What are Useful Characteristics of LWs?

To finalize our analysis, we want to stress the key characteristics of LWs that the model underlines. First, LWs are the instrument that discloses to creditors the restrictions on X contained in ω. Second, they also provide them with information about their payoffs in each possible resolution method by describing the value of the firm. In the model, these two characteristics of LWs are necessary for debtholders

[10] See bankruptcy reform proposals by Jackson (2014).

to correctly estimate the probability of repayment and to price the debt accordingly. That is, creditors need to observe X to tame moral hazard by making the price of debt sensitive to firm choices. More generally, an important lesson may apply to LWs outside of our model:

Implication 5 *In order for LWs to lessen the time inconsistency problem, we need creditors to have access to the resolution plan and understand the implied priority order of payments under each possible resolution method.*

When thinking about this characteristic of "transparency" of the LWs, it is interesting to consider the scenario in which regulators have commitment not to intervene. In this scenario, the simple disclosure of X would imply that creditors would themselves impose limits on the firm's choices (through unregulated LWs) by having the price of debt depend on X. As Implication 2 stated, if externalities are not important, this may be an efficient equilibrium without regulation and plausibly a desirable situation for society. An important question, however, is whether with unregulated LWs such a level of disclosure of the firm's choices would be feasible, or if it would be incompatible with competition among different financial firms. If the latter were the case, there may be a role for regulation of LWs (imposing limits on X without disclosing the X itself) even in a world with perfect commitment to not bail out.

5. CONCLUSION

It has been argued that living wills are a promising new tool in supervision that may help policymakers alleviate the TBTF problem. In this article we have described the mechanism through which they may achieve that objective and under what conditions they are more likely to succeed. A key insight from our analysis is that the requirement for financial firms to file living wills is not equivalent to regulators tying their hands ex ante so that they are not able to intervene with a bailout in the event of financial distress. Instead, the requirement that firms have living wills in place is meant to make the outcomes from bankruptcy better for society. This has two beneficial effects. First, it directly lessens the moral hazard problem that the possibility of bailouts creates by expanding the set of choices of the firm on which supervisory requirements are imposed: If the size or the complexity the firm has chosen is such that the plan for resolution unveils important costs to using bankruptcy, regulators can require the firm to adjust its structure. This makes the ex-post choice of policymakers not to intervene more attractive and hence more likely. If unassisted failures are more likely, in turn, a second, indirect effect arises: Debtholders who stand to lose in those failures will increase their monitoring of hard-to-regulate risk choices of the firm, again decreasing the moral hazard problem.

Because regulators are not ruling out ad-hoc funding, in particular cases when intervention seems a better choice (such as when a firm fails in the midst of adverse aggregate conditions), bailouts are still an available tool for policymakers. However, through their increased monitoring of firm characteristics key to the strength of financial institutions, living wills may not only decrease the probability of intervention, but also the size of the public funds involved in interventions when they happen, saving society the costs these transfers may involve.

It is important to keep in mind, however, that there are tradeoffs involved in the requirement of living wills: There are costs to the firms of writing a living will and making the changes to its structure that regulators deem necessary to make their unassisted resolution through bankruptcy viable. Hence, regulators need to exercise care that (i) the company does invest the necessary resources to produce a truthful and useful living will, and (ii) that the costs of the changes and the resources necessary to craft a good living will do not wipe out the expected benefits to society from having more resolvable and resilient financial firms.

It is important that the implications from our analysis are taken with caution and understood within the context of our assumptions. For simplicity of the analysis, we have assumed that the policymaker is able to perfectly evaluate the implications (for stakeholders of the firm, as well as for the economy as a whole) of resolving a financial firm in distress of certain characteristics through bankruptcy or the Orderly Liquidation Authority, and we compare that scenario with that of a bailout or a market-based resolution. Further research that considers more realistic informational constraints on the policymakers regulating the living will review process is likely to qualify our implications.

Moreover, our analysis highlights the important role that externalities play in determining the efficiency of an institution such as the Orderly Liquidation Authority, or even the regulation of living wills. Our framework allows us to compare welfare with regulated living wills in place to welfare without them, in a world where we rely on the debt market to monitor firm choices and to price debt accordingly. The importance of externalities is a key parameter in this comparison. Quantitative explorations of the advantages and disadvantages of different methods for resolution, including quantification of potential externalities, would greatly enhance our understanding of the potential for living wills and any related proposals to end the TBTF problem.

REFERENCES

Acharya, Viral V., Thomas F. Cooley, Matthew P. Richardson, and Ingo Walter. 2010. *Regulating Wall Street: The Dodd-Frank Act and the New Architecture of Global Finance*. Hoboken, N.J.: John Wiley & Sons, Inc.

DeYoung, Robert, Michal Kowalik, and Jack Reidhill. 2013. "A Theory of Bank Resolution: Technological Change and Political Economics." *Journal of Financial Stability* 9 (December): 612–27.

Grochulski, Borys. 2011. "Financial Firm Resolution Policy as a Time-Consistency Problem." Federal Reserve Bank of Richmond *Economic Quarterly* 97 (Second Quarter): 133–52.

Gruenberg, Martin. May 12, 2015. "A Progress Report on the Resolution of Systemically Important Financial Institutions." Speech to the Peterson Institute for International Economics, Washington, D.C. Available at https://www.fdic.gov/news/news/speeches/spmay1215.html.

Herring, Richard J., and Charles W. Calomiris. 2011. "Why and How to Design a Contingent Convertible Debt Requirement." Wharton Financial Institutions Center Working Paper 11-41 (April). Available at http://dx.doi.org/10.2139/ssrn.1815406.

Jackson, Thomas H. July 9, 2014. "Building on Bankruptcy: A Revised Chapter 14 Proposal for the Recapitalization, Reorganization, or Liquidation of Large Financial Institutions." Manuscript, Stanford University.

Jarque, Arantxa, and David A. Price. 2015. "Living Wills: A Tool for Curbing 'Too Big to Fail.'" Federal Reserve Bank of Richmond *2014 Annual Report,* 5–17.

Pellerin, Sabrina R., and John R. Walter. 2012. "Orderly Liquidation Authority as an Alternative to Bankruptcy." Federal Reserve Bank of Richmond *Economic Quarterly* 98 (First Quarter): 1–31.

White, Phoebe, and Tanju Yorulmazer. 2014. "Bank Resolution Concepts, Trade-offs, and Changes in Practices." Federal Reserve Bank of New York *Economic Policy Review* 20 (December): 1–21.

Economic Quarterly—Volume 101, Number 3—Third Quarter 2015—Pages 225–244

What We Know About Wage Adjustment During the 2007–09 Recession and Its Aftermath

Marianna Kudlyak

In standard economic models, unemployment and real wage growth are tightly connected, moving at nearly the same time in opposite directions (Daly and Hobijn 2014). Historically, however, the relationship between aggregate real wage growth and unemployment has been weak and, specifically, aggregate wage growth has remained flat during the 2007–09 recession and its aftermath while unemployment has exhibited substantial swings (Figure 1). This experience has led many observers, including some policymakers, to question whether the low real aggregate wage growth during the current recovery indicates a weak labor market beyond what is measured by the official unemployment rate.

The category "aggregate wage" summarizes wages of all employed workers, and thus aggregate wage growth reflects actual changes of workers' wages, changes in the composition of workers, and changes in the composition of jobs. Some of these changes are related to underlying structural trends in the economy while others constitute the economy's response to the business cycle shocks and are more indicative of cyclical resource utilization in the labor market. Consequently, it is important to look beyond the aggregate statistics to understand

Kudlyak is an economist at the Federal Reserve Bank of Richmond. The author thanks Miki Doan, Robert Hetzel, Pierre-Daniel Sarte, and John Weinberg for useful comments and suggestions. The views expressed here are those of the author and do not reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: marianna.kudlyak@rich.frb.org.

Figure 1 Aggregate Real Wage Growth and Unemployment Rate

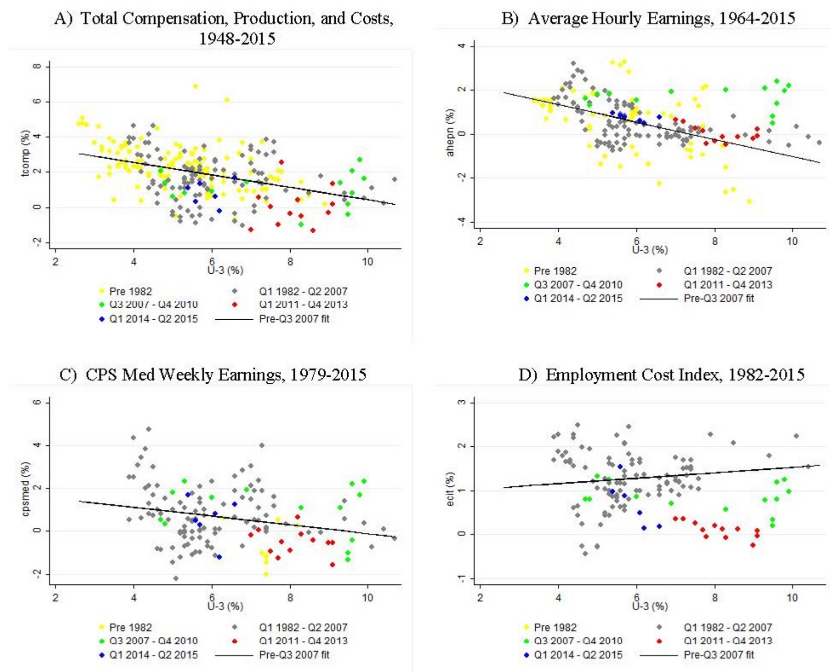

Notes: Each dot represents a quarterly observation on the unemployment rate and annual real wage growth. Hourly compensation in the nonfarm business sector is from the Bureau of Labor Statistics' Production and Costs tabulations. Average hourly earnings of production and non-supervisory workers are from the Current Employment Statistics. The employment cost index is from the National Compensation Survey. Median weekly earnings are from the Current Population Survey.

the behavior of real wages and its relation to the health of the labor market in the aftermath of the 2007–09 recession.

In this article, we review recent literature that studies the changes in the components of the aggregate wage over time and, specifically, after the 2007–09 recession. The review focuses on disaggregating the aggregate real wage growth into its changes due to: wage changes of workers who switch from one job to another, wage changes of workers who remain employed at the same job, wage changes of workers who move in and out of employment, and changes in respective shares of these groups. In principle, the change in the aggregate real wage from one period to another can be disaggregated along a few alternative dimensions—for example, the changes in the industry (or occupational) wages and industrial (or occupational) composition of employment, or the changes in wages of workers by sociodemographic characteristics and the educational and experience composition of the workforce. The decomposition in this article is motivated by the earlier (for example, Barro 1980) and recent macrolabor literatures that converge on a two-part argument regarding wage changes that influence job creation decisions of firms. First, wages of workers who do not change jobs do not directly influence the firms' job creation decisions. Instead, firms take into account the entire expected present value of wages to be paid to a newly hired worker (Haefke, Sonntag, and van Rens 2013; Kudlyak 2014; and Basu and House 2015). Second, wages of newly hired workers change much more over the business cycle than wages of workers who remain at the same job (see, for example, Bils 1985; Pissarides 2009; and Martins, Solon, and Thomas 2012), and thus movements in average wages are not the relevant statistics that influence firms' job creation decisions. Having disaggregated aggregate wage growth along the dimensions described above, we further examine the changes in the composition of each dimension along high- versus low-paying jobs and industrial or occupational makeup whenever the relevant studies are available.

Reviewing the literature, we find that wage changes of workers employed from one period to another are procyclical and the majority of the procyclicality is due to the wage changes of workers who change jobs from one period to another. The compositional changes in the aggregate wage are typically countercyclical. Consequently, aggregate real wage growth is procyclical but not in a statistically significant sense (Daly, Hobijn, and Wiles 2012). The review of the literature reveals that the aggregate wage growth in the aftermath of the 2007–09 recession is relatively low because of the relatively low contribution of the typically procyclical wage growth of job changers and because of the relatively large contribution of the typically countercyclical composition effect associated with transitions from part-time to full-time employment (Daly, Hobijn, and Wiles 2012). Despite the fact that the overall contribution of job changers to aggregate wage growth is relatively low in this recovery, the wage gains of workers who do switch jobs are high (Mustre-del-Rio 2014). In light of the recent literature on declining business dynamism (see, for example, Davis and Haltiwanger 2014), it may be that the relatively low rate of job-to-job switches is structural in nature.

Given the centrality of the question of wage growth for economic recovery, there have been a number of works studying wage growth in the aftermath of the 2007–09 recession.[1] This article's intent is to provide a roadmap for readers of recent empirical works on wage dynamics in the context of wage changes as it relates to job creation decisions of the firms.[2] The review is descriptive and many questions remain as to sources behind the weak growth of the components of the aggregate wage growth.

The article is structured as follows. Section 1 presents a framework for decomposition of aggregate wage growth into wage changes of workers who do not change jobs, wage changes of workers who switch jobs, wage changes of workers who transition between nonemployment and employment, and the changes in the respective shares of these groups. Section 2 summarizes Daly, Hobijn, and Wiles (2012), whose findings speak to the multiple components of the decomposition. Section 3 discusses evidence on wage adjustment of workers who do not change jobs. Section 4 discusses evidence on wage adjustment of workers who switch jobs. Section 5 discusses some evidence on wage adjustment of workers who transition from nonemployment to employment. Section 6 discusses the types of jobs workers transitioned to during the recovery from the 2007–09 recession. Finally, Section 7 concludes.

1. A FRAMEWORK FOR DECOMPOSITION OF AGGREGATE WAGE GROWTH

The change in real aggregate average wages between $t-1$ and t can be decomposed such that it explicitly distinguishes between changes in the wages of job switchers and job stayers, relative wages of workers who enter and exit the workforce, and changes in the shares of each group of workers. Let W_t denote the aggregate average wage in period t. Let C denote the set of workers who remain at the same job from period $t-1$ to t. Let S denote the set of workers who work in both periods but switch jobs between $t-1$ and t. Let X denote the set of workers who exit employment between $t-1$ and t, i.e., those who are employed in $t-1$ and separated from employment in period t. Finally, let N denote the set of workers who enter employment between $t-1$ and t, i.e., those who are unemployed or out of the labor force in period t and employed in period $t-1$. Let s_{it} denote the share of workers in set i of the workforce in period t. Let w_{it} denote the real wage of the

[1] As this article is being written, new works are published on the topic, and thus the review is necessarily incomplete. See, for example, Barrow and Faberman (2015).

[2] An interested reader is referred to a review in Kudlyak (2009) and a chapter in the forthcoming *Handbook of Macroeconomics* by Basu and Hause (2015) for the theoretical background.

workers in set i in period t. Then the aggregate wage change can be decomposed as follows:[3]

$$W_t - W_{-1} = \sum_{i \in C,S,N} s_{it} w_{it} - \sum_{i \in C,S,X} s_{it-1} w_{it-1}$$

$$= \sum_{i \in C} s_{it-1} \left(w_{it} - w_{it-1} \right) \text{ (1) wage growth effect of stayers}$$

$$+ \sum_{i \in S} s_{it-1} \left(w_{it} - w_{it-1} \right) \text{ (2) wage growth effect of changers}$$

$$+ \sum_{i \in N} s_{it} \left(w_{it} - W_{t-1} \right) \text{ (3) composition effect of the entrants to employment}$$

$$- \sum_{i \in X} s_{i-1t} \left(w_{it-1} - W_{t-1} \right) \text{ (4) composition effect of the exiters from employment}$$

$$+ \sum_{i \in C,S} \left(w_{it-1} - W_{t-1} \right) \left(s_{it} - s_{it-1} \right) \text{ (5) composition effect of stayers and switchers}$$

$$+ \sum_{i \in C,S} \left(w_{it} - w_{it-1} \right) \left(s_{it} - s_{it-1} \right) \text{ (6) composition effect from cross-term.} \quad (1)$$

The first two terms on the right constitute the contribution to the change in the aggregate real wage from the changes of wages of job stayers and job changers, respectively, holding the shares of these groups in the pool of all employed constant at their period $t-1$ levels. The third term captures the effect of new entrants into employment: It

[3] To obtain the decomposition, add and subtract $W_{t-1} \sum_{i \in C,S} \left(s_{it} - s_{it-1} \right) s_{it-1}$, $\sum_{i \in C,S} s_{it-1} w_{it}$, $\sum_{i \in C,S} w_{it-1} \left(s_{it} - s_{it-1} \right)$ from the left-hand side and note that $W_{t-1} \sum_{i \in C,S} \left(s_{it} - s_{it-1} \right) s_{it-1} = W_{t-1} \left(\sum_{i \in N} s_{it} w_{it} - \sum_{i \in X} s_{it-1} w_{it-1} \right)$. After rearranging, equation (1) follows. Note that $s_{Ct} + s_{St} + s_{Nt} = 1$.

captures the new entrants' wage relative to the initial aggregate wage weighted by their share in period t. The fourth term captures the effect of workers who exit employment: It captures the exiting workers' wages relative to the initial aggregate average wage weighted by their share in period t. Daly, Hobijn, and Wiles (2012), for example, find that the workers who transition from employment to unemployment usually come from the low end of the wage distribution and workers who enter employment from nonemployment usually enter at even lower wages. The fifth term reflects the changing shares of job stayers and job changers, weighted by deviation of their initial wages from the initial aggregate wage. That is, if the shares shifted toward the groups with higher wages than the initial average, such a shift would represent a positive contribution to the change in the aggregate real wage. Finally, the sixth term captures the remaining cross-term.

Movements in any of these components affect movements in the aggregate wage. In particular, a large literature documents that real wages of job stayers are weakly procyclical and the wages of job changers are highly procyclical (Pissarides 2009).[4] Consequently, effects (1) and (2) should lead to procyclical aggregate wage changes. However, layoffs and hiring are not randomly distributed over the business cycle: Low-wage workers are more likely to be displaced in recessions and thus a measure of the aggregate wage gives more weight to low-skill workers in expansions than in contractions (Solon, Barsky, and Parker 1994). These composition effects might counteract the procyclicality of the wage changes of job stayers or job changers.

Note that wage changes of, for example, job changers can be due to the "true" wage change of workers who switch jobs (i.e., workers receiving higher pay for the same type of job) and/or due to a change in the types of jobs the workers switch to. The same applies to the effects of job stayers or the effects of workers transitioning from nonemployment to employment or vice versa. In particular, McLaughlin and Bils (2001) and Moscarini and Postel-Vinay (2014), among others, document substantial cyclical upgrading; that is, during economic recoveries workers switch to better-paying industries, occupations, or firms. To our knowledge, however, the evidence that explicitly distinguishes between such job-related compositional effects is scarce.

The decomposition above is analogous to the shift-share decomposition for productivity growth presented in Foster, Haltiwanger, and Krizan (2001) and is a simplified version of the aggregate *median* wage

[4] A relative weak procyclicality of job stayers' wages can be rationalized by efficiency wage models, implicit contracts, or insider-outsider models.

decomposition developed by Daly, Hobijn, and Wiles (2012), which is discussed in Section 2.

The next sections discuss recent empirical evidence on wage adjustment in the 2007–09 recession and its recovery in the context of the decomposition presented in equation (1).

2. EMPIRICAL EVIDENCE ON THE WAGE GROWTH EFFECT AND THE COMPOSITION EFFECT FROM DALY, HOBIJN, AND WILES (2012)

Daly, Hobijn, and Wiles (2012) present a detailed decomposition of aggregate wage adjustment during the 2007–09 recession and its aftermath along the lines presented in equation (1). They focus on the decomposition of the changes in the aggregate median wage growth of full-time workers. The focus on the wage changes of full-time workers as opposed to the wages of all employed workers allows them to explicitly consider the contribution of part-time workers to the composition effects (3) and (4) described in equation (1), in addition to the contributions of the unemployed or workers out of the labor force. Such effects might be especially important in the aftermath of the 2007–09 recession. First, involuntary part-time employment in the aftermath of 2007–09 was higher than during previous economic recoveries and thus has been one of the topics at the center of the policy discussions regarding labor market resource underutilization. Second, involuntary part-time workers typically receive lower wages than full-time employed workers even after controlling for observable characteristics and broad occupational and industry categories (see, for example, Canon et al. 2015). In addition, Daly, Hobijn, and Wiles (2012) use a novel continuous-distribution version of the shift-share analysis presented in equation (1) to decompose the growth in the aggregate *median* wage into the changes in individual components versus the growth in aggregate average wage.[5] As noted in Daly, Hobijn, and Wiles (2012), as opposed to the average wage, median is not affected by fluctuations in overtime, overtime pay, or trends in the average work week for full-time workers.

Daly, Hobijn, and Wiles (2012) use the Current Population Survey (CPS) microdata, 1980–2011, and construct median usual weekly earnings (MWE) of full-time wage and salary workers. The Bureau of Labor Statistics (BLS) publishes MWE of full-time wage and salary

[5] The intuition behind the Daly, Hobijn, and Wiles (2012) decomposition is similar to the intuition described by equation (1). The decomposition functions like a shift-share analysis but works with distribution functions rather than means. An interested reader is referred to the original article for the exact expressions.

workers on a quarterly basis, and it is the only wage series for which the underlying microdata are publicly available.[6] The microdata for the MWE come from the CPS outgoing rotation groups files (i.e., the data collected from the individuals in their fourth or eighth month of the CPS interview). As such, the published four-quarter aggregate wage growth series is calculated from wages of workers in their fourth or eighth month of the CPS interview in period t and wages of workers who are in their fourth or eighth month of the CPS interview in period $t-1$, i.e., two different samples of workers. Instead, Daly, Hobijn, and Wiles construct the *matched* MWE from the sample of individuals who are present in both periods—in their interview month eight in period t and in their interview month four in period $t-1$. Such an approach allows linking the changes in aggregate wage growth to the changes in its components, sidestepping the questions of individuals being in and out of the survey due to the survey design.[7]

Daly, Hobijn, and Wiles (2012) observe that workers employed full time in both periods account for, on average, about 90 percent of all full-time wage earners. Consequently, these workers' wages are likely to drive most of the movements in aggregate wages. In contrast, workers exiting from (entering into) full-time employment make up only a small share of overall wage earners. For example, those exiting from full-time employment to unemployment and entering from unemployment to full-time employment make up only 2.7 percent and 2.6 percent, respectively. However, the flows into and out of full-time employment are typically drawn from below the median wage computed across all employed workers and thus can generate a quantitative impact on the aggregate median wage. For example, 72.8 percent of those who enter from unemployment and 80 percent of those who enter from part-time employment or out of the labor force enter at or below the median wage.

Daly, Hobijn, and Wiles (2012) find that the wage growth effect is strongly procyclical (the coefficient on unemployment is -0.222) and accounts for the majority of the variance of real wage growth over time (91.2 percent). In contrast, the composition effect is countercyclical (the coefficient on unemployment is 0.110). Consequently, real median wage growth is procyclical but not in a statistically significant sense

[6] To understand the compatibility of the aggregate wage series from the CPS to other widely used series, Daly, Hobijn, and Wiles (2012) show that the MWE compares well with average hourly earnings of production and non-supervisory workers in the private sector, compensation per hour in the nonfarm business sector, and the employment cost index. In particular, the MWE exhibits similar coincident movements with these series, and the correlation with the other series is at least 0.60.

[7] The correlation between matched MWE and MWE is 0.79.

(the coefficient on unemployment is −0.112). The authors find that the relative importance of the wage growth effect and the composition effect varies considerably over the business cycle. In downturns, the composition effect rises in importance, offsetting more than half of the procyclicality of the wage growth effect.

Examining the effects of job stayers and job changers separately, Daly, Hobijn, and Wiles (2012) find that the effect of workers who change jobs on the aggregate wage is larger than the effect of workers who remain in the same job from period to period. When labor markets are tight, the job changers' effect is amplified by an increase in the share of the job changers and an increase in the fraction of them moving from below to above the median wage. In labor market downturns, the effects of job changers and job stayers converge as both the share of job changers falls and their relative earnings gains subside.

Interestingly, the unemployment margin is relatively unimportant for the variance and cyclicality of aggregate median full-time wage growth. The net composition effect of unemployment is a sum of the composition effect of workers who transition from employment to unemployment and the composition effect of workers who transition from unemployment to employment. Daly, Hobijn, and Wiles (2012) note the "unemployment penalty," whereby the unemployed workers are typically re-employed at even lower wages than their pre-unemployment wage (due to loss of human capital or some other effects). If the shares of those who enter and exit employment from/to unemployment are similar, the effect of the unemployment component on the aggregate wage is negative. However, because the share of exits from full-time employment to unemployment increases in contractions, the negative effect is somewhat counteracted during downturns. Consequently, the authors find that the net composition effect of unemployment on aggregate wage change is small, negative, and only weakly countercyclical.

Daly, Hobijn, and Wiles (2012) find that much of the composition effect comes through the part-time employment margin. The magnitude of the part-time and self-employment effect relative to other margins owes to the fact that a larger fraction of flows into and out of part time occur below the aggregate medium wage. In contrast to the unemployment margin, the part-time effect does not have offsetting components: In downturns, the share of exits from full time to part time rises and there is little change in the earnings difference between entrants and exits between full-time employment and part-time employment.

Finally, the Daly, Hobijn, and Wiles (2012) results demonstrate that the combined wage growth effect of job changers and job stayers in the aftermath of the 2007–09 recession is lower than after the previous downturns. In addition, there has been a larger than usual countercyclical contribution of the composition effect. Much of the countercyclical contribution of the composition effect after 2010 is due to part-time employment (Daly, Hobijn, and Wiles 2012, Figure 7).

3. ADDITIONAL EVIDENCE ON CHANGES OF WAGES OF JOB STAYERS

A series of recent articles provides evidence on the change of wages of job stayers. These studies focus on nominal wage rigidities; nevertheless, it is instructive to review the main evidence and the accompanying arguments.

Daly and Hobijn (2015) argue that downward nominal wage rigidity is an explanation for the lack of movement in the opposite direction of wages and unemployment during the 2007–09 downturn and the subsequent recovery. That is, the aggregate wage did not adjust downward during the 2007–09 recession due to nominal wages being more rigid than during previous downturns combined with low inflation. Consequently, they argue that wage growth has not picked up during the economic recovery because firms need to work off a stockpile of pent-up wage cuts.

In contrast, Elsby, Shin, and Solon (2014) argue that while the data show downward nominal wage rigidities, the slight increase in the rigidities during the 2007–09 recession is a part of a pre-existing trend. Similar to Daly, Hobijn, and Lucking (2012), Elsby, Shin, and Solon (2014) note a possible role for the low inflationary environment in the reduced procyclicality of real wages during the 2007–09 recession. However, Elsby, Shin, and Solon (2014) emphasize that wages of job stayers do not play a direct allocative role for employment. Below, we review these two studies in some more detail. Then we provide a brief description of recent evidence on countercyclicality of work effort of job stayers that might render the effective wage of job stayers procyclical even though the actual paid wage does not change over the business cycle.

Downward Nominal Wage Rigidities

Direct evidence of nominal wage rigidities is presented in Daly, Hobijn, and Lucking (2012). They use the CPS microdata from 1980 to 2011 (hourly or salaried workers who did not change jobs from year to year) and show that nominal wage rigidity appears to increase during recessions and lag during recovery. During the 2007–09 downturn, the share of workers who report no nominal wage change increased from 11.2 percent to 16 percent, the highest level observed in any period starting from 1980. Nominal wage rigidities during 2007–09 were prevalent among all educational groups, while during the 1990–91 recession, the higher-educated group (college) did not experience an increase in nominal wage rigidity. The finance, construction, and manufacturing industries experienced the run-up in nominal wage rigidity. Daly, Hobijn, and Lucking (2012) conclude that a somewhat higher rate of inflation would grease the wheels of the labor market by allowing real wages to fall.[8]

Daly and Hobijn (2015) argue that wage growth during the 2007–09 recovery is slow because many firms were unable to reduce wages during the recession and they must now work off a stockpile of pent-up wage cuts. To support the argument, using cross-industry data, they show that industries that were least able to cut wages during the downturn and therefore accrued the most pent-up cuts have experienced relatively slower wage growth during the recovery.[9]

Evidence against Unusually High Nominal Wage Rigidities During the 2007–09 Recession

Elsby, Shin, and Solon (2014) use a different CPS subsample and a slightly different data treatment than the one used in Daly, Hobijn, and Lucking (2012) and argue that the nominal wage rigidities were not unusually high during the 2007–09 recession and that the slight increase is part of a long-run trend.

Using tenure supplements to the CPS, 1979–2012, Elsby, Shin, and Solon (2014) separately examine nominal wage changes for hourly and non-hourly workers; for non-hourly workers, they use usual weekly earnings instead of constructing the hourly wage by dividing weekly earnings by weekly hours as in Daly, Hobijn, and Lucking (2012).

[8] The argument is based on Akerlof, Dickens, and Perry (1996).

[9] The discussion of the relationship between wage inflation and price inflation is beyond the scope of this review. An interested reader is referred to, for example, Daly and Hobijn (2014), who argue that nominal wage rigidity affected the aggregate relationship between the unemployment rate and wage growth during the past three recessions and recoveries and has been especially pronounced during and after the 2007–09 recession. In particular, they present a model of monetary policy with downward nominal rigidities and show that both the slope and curvature of the Phillips curve depend on the level of inflation and the extent of downward nominal wage rigidities. They show that downward wage rigidity results in the "bending" of the Phillips curve whereby the wage growth during a recovery is lower than it is at the same level of the unemployment rate when unemployment is rising, i.e., during the downturn.

Using the real wage levels, Elsby, Shin, and Solon (2014) find that wages are procyclical.[10] However, during the 2007–09 recession, men's real wages did not fall, while women's real wages adjusted downward. Using data on year-to-year nominal wage changes of job stayers, they report observing "both a substantial minority of workers reporting the same nominal wage in adjacent years (suggesting nominal wage stickiness), but also a substantial minority reporting nominal wage cuts (suggesting nominal wage flexibility). In addition, recent data spanning the Great Recession suggest only a modest rise in the incidence of nominal wage freezes.[11] The authors conclude that wages are nominally rigid, and that there is an upward trend in nominal wage rigidity.

Similar to Daly, Hobijn, and Lucking (2012), Elsby, Shin, and Solon (2014) note a possible role for the inflationary environment in reducing the procyclicality of men's real wages during the 2007–09 recession. In particular, they write that "[a]t the outset of the recession of the early 1980s, inflation was unusually high, and employers could reduce real wages substantially even while granting nominal wage *increases*. This was still somewhat true in the recession of the early 1990s, when annual inflation was about 4%. But during the Great Recession, especially in 2009, the inflation rate was lower, and substantial real wage cuts would have required nominal wage cuts."

Countercyclical Work Effort

Additional evidence on the behavior of wages of workers who do not change jobs comes from the recent evidence on countercyclical work effort. If employed workers exert higher effort in recessions, then even if their wage does not change from period to period, the workers' effective wage is lower because they are paid less for a unit of output. Under such a scenario, even though actual wages of job stayers do not change, the effective wages decrease in downturns and increase during recoveries.[12]

[10] Elsby, Shin, and Solon (2014) find that regression-adjusted analysis or accounting for unobserved heterogeneity reveals more procyclicality.

[11] Elsby, Shin, and Solon (2014) state that the documented patterns are consistent with those documented by Daly, Hobijn, and Lucking (2012) with two differences: "Daly et al. (2012) divide reported weekly earnings by reported weekly hours to get their nominal wage measure for non-hourly workers. As expected, this leads to considerably smaller spikes at zero nominal wage change. Second, for hourly workers, Daly et al. estimate a substantial dip in the frequency of zero nominal wage change in the years preceding the Great Recession. In contrast, our estimates in Table 5 do not show a drop-off after 2003–04."

[12] See Bils, Chang, and Kim (2013) on how rigid wages of job stayers can amplify the job creation over the business cycle if work effort varies.

The argument about the countercyclicality of work effort is related to a large earlier literature on "efficiency wages" whereby high unemployment induces higher work effort of the employed workers and thus firms do not lower wages of incumbent workers during recessions (Stiglitz 1974). The earlier evidence on efficiency wages is primarily nondirect. Direct evidence on the cyclicality of effective wages is scarce due to the lack of measures of individual workers' output. The exception is Anger (2011), who finds that effective wages are more procyclical than actual wages using German data.

A recent direct piece of evidence of higher worker effort in recessions comes from Lazear, Shaw, and Stanton (2013). Using data spanning June 2006 to May 2010 on individual worker productivity from a large firm, they find evidence that increased productivity during the downturn was mostly due to increases in workers' effort rather than the improved composition of the workforce. That is, firms get more effort from fewer workers, i.e., "making do with less."

The literature on "making do with less" during recessions appears to contradict the earlier literature on labor hoarding (Oi 1962), whereby firms resist firing workers during downturns and instead keep partially or entirely idle workers on their payrolls. To the extent that there have been trend-related decreases in the costs of hiring and training (for example, due to faster depreciation of skills during non-activity) so that labor hoarding has become less prevalent, higher effort during recession can rationalize the increased incidence of downward nominal or real wage rigidity.

4. RECENT EVIDENCE ON WAGES OF JOB-TO-JOB CHANGERS

Direct evidence on wages of job changers is presented by Mustre-del-Rio (2014). Using data on nominal wages of hourly workers in the private sector from the Survey of Income and Program Participation, Mustre-del-Rio constructs the series of monthly wage growth for job switchers (without an intervening spell of joblessness). During 1998–2012, the contemporaneous correlation of wage growth of job switchers with quits (from the Job Openings and Labor Turnover Survey) is 0.50. The correlation increases to 0.73 in the post-2007 period, suggesting an even stronger relationship between wage changes of job switchers and quits.

A similar conclusion is derived from the Automatic Data Processing's Workforce Volatility Index. Mustre-del-Rio (2014) finds that job changers' wage growth rose from 4.3 percent per quarter in Q1:2013 to 5.6 percent in Q3:2014. Job changers in leisure and hospitality reached an average wage growth of 7.7 percent, switchers in professional and business services reached 6.8 percent, switchers in the education and health sector observed gains of 3.9 percent, and switchers in the manufacturing sector reached 4.2 percent.

How can one reconcile evidence presented by Mustre-del-Rio (2014) that the correlation between wages of job switchers and job quits is even higher after 2007 than it was historically with evidence from Daly, Hobijn, and Wiles (2012) that job changers' wage growth contributed relatively less to the aggregate wage growth during the recovery from the 2007–09 recession? As noted by others, during the recovery from the 2007–09 recession, job-to-job transitions failed to pick up at their historical pace. Higher pay for the same job elsewhere or better jobs available elsewhere enhance the growth of job changers' wages. But if there are no better jobs available, there are no quits. Moscarini and Postel-Vinay (2014) argue that the job ladder "failed" after the 2007–09 recession. They provide evidence that job-to-job quits collapsed, especially toward large, high-paying employers. Moscarini and Postel-Vinay argue that new jobs at small employers, the traditional first step of the ladder out of unemployment, vanished, in part because large employers stopped poaching and did not create room for hiring at the bottom. Section 6 reviews additional studies on the types of jobs created during the recovery from the 2007–09 recession, but the evidence is mixed.

5. WAGES OF HIRES FROM NONEMPLOYMENT

Recent evidence on wages of workers who transition from nonemployment to employment comes from the studies of wages of recent college graduates. Hobijn and Bengali (2014) argue that the median starting wages of recent college graduates declined to a greater degree than the wages of all full-time workers and that this effect cannot be attributed to the composition of jobs in which the recent graduates are employed. They interpret the evidence as support for the hypothesis of a relatively weak recovery.

In particular, Hobijn and Bengali argue that median nominal wages of recent college graduates (21–25-year-old workers with college degrees; not necessarily newly hired) have not kept pace with median earnings for all workers over the past six years. The wage gap in the current recovery is substantially larger and has lasted longer than the previous recovery.

Hobijn and Bengali then examine whether these developments are due to a composition effect. They present evidence that the distribution of recent college graduates across occupations has roughly remained the same across 2007, 2011, and 2014, while the cumulative nominal wage growth of the graduates between 2007 and 2014 has lagged that of all full-time workers. Based on this evidence, Hobijn and Bengali argue that the larger gap between wages of all full-time workers and the wages of recent graduates represents slow growth in starting salaries for graduates rather than a shift in types of jobs.

Abel, Deitz, and Su (2014) analyze the types of jobs held by recent college graduates in greater detail. They argue that recent graduates are typically "underemployed," i.e., working in jobs that typically do not require a bachelor's degree,[13] and find that this phenomenon is not specific to the 2007–09 recession as graduates always have a harder time finding well-paid jobs during recessions (Kahn 2010). Abel, Deitz, and Su (2014), however, find that the quality of the jobs held by the underemployed graduates has declined after the 2007–09 recession to a greater degree than in the past. Consequently, Abel, Deitz, and Su (2014) suggest that low wages of new college graduates are likely due to a composition effect of the types of jobs in which they are employed.

Abel, Deitz, and Su (2014) define a college graduate as underemployed if the occupation she is working in requires a bachelor's degree less than 50 percent of the time, as defined by O*NET Education and Training Questionnaire.[14] The underemployment rate is then the ratio of the number of underemployed college graduates to all employed college graduates. They find that younger college graduates are always more underemployed than older college graduates; however, they were underemployed to a larger degree during the 2007–09 recession. The authors speculate that while there appears to be a cyclical component to underemployment among recent college graduates, the broader V-shaped pattern in the underemployment rate over the past two decades is also consistent with new research arguing that there has been a reversal in the demand for cognitive skills since 2000.[15]

[13] Their definition of recent college graduates is slightly different from the one by Hobijn and Bengali (2014) as they include 22–27-year-old, not-in-school workers.

[14] They use the answer to the question "If someone were being hired to perform this job, indicate the level of education that would be required?" and consider college education to be a requirement for a given occupation if at least 50 percent of the respondents working in that occupation indicated that a bachelor's degree is necessary to perform the job.

[15] Abel, Deitz, and Su (2014) write, "[a]ccording to this research, businesses ramped up their hiring of college-educated workers in an effort to adapt to the technological changes occurring during the 1980s and 1990s. However, as the information technology revolution reached maturity, demand for cognitive skill fell accordingly. As a result, during the first decade of the 2000s, many college graduates were forced to move down the occupational hierarchy to take jobs typically performed by lower-skilled workers. From this perspective, the relatively low underemployment rates among recent college graduates at the peak of the technology boom around 2000 may in fact be an outlier, while the recent rise in underemployment represents a return to more typical conditions."

Delving deeper, the authors show that among the underemployed recent college graduates, the share of those with well-paid non-college jobs is decreasing while the share of those with low-paid non-college jobs is increasing. The authors show that these developments have been taking place since 2001.

Importantly, Abel, Deitz, and Su (2014) show that unemployment and underemployment differ markedly across majors. Specifically, majors providing technical training (quantitative and analytical skills) had the highest shares of graduates working in jobs that require a degree (engineering and math and computers majors). In addition, education and health majors also had high shares of workers employed in jobs that require a college degree, likely reflecting the growth of these sectors in recent years.

6. TYPES OF JOBS

Mester and Sen (2013) examine employment growth, distinguishing between qualities of jobs. In particular, they examine employment growth within higher-than-average and lower-than-average wage industries during 1990–2012. They find that in the 1991 and 2001 recessions, when employment begins to grow in a recovery, the first jobs added are typically in industries that are relatively low-paying, while higher-paying jobs are added later as the economy and employment continue to expand. However, after the 2007–09 recession, the story is different—instead of losing higher-paying jobs at a faster pace than lower-paying jobs during the recession and recovering lower-paying jobs more quickly at the start of the recovery, higher-paying jobs were cut less sharply during the 2007–09 recession and have been added at the same or a faster pace than lower-paying jobs during the recovery.

Mester and Sen note, however, that the analysis is not without caveats. Their classification of industries is not fixed through the sample period. Instead, for each month, they classify industries into higher-paying and lower-paying according to whether their average hourly earnings are above or below, respectively, the national average for all private industries in that month. Consequently, some industries switch from one category to another during the sample period. Specifically, Mester and Sen note that, in 2006, manufacturing switched from being a high-paying industry to a low-paying industry. Repeating that analysis and classifying manufacturing as a high-paying industry after 2006, Mester and Sen find that the pattern from the previous recessions persists into the 2007–09 one—employment growth in high-paying jobs picks up later in the recovery as compared to employment growth in low-paying jobs. As such, evidence on the changes in the quality of jobs after the 2007–09 recession is mixed.

7. CONCLUSIONS

The review of the literature above reveals that the aggregate wage growth in the aftermath of the 2007–09 recession is low because of the relatively low contribution of the typically procyclical wage growth of job switchers and because of the relatively large contribution of the typically countercyclical composition effect associated with transitions from part-time to full-time employment (Daly, Hobijn, and Wiles 2012). Despite the fact that the overall contribution of job switchers to aggregate wage growth is relatively low in this recovery, the wage gains of workers who do switch jobs are high (Mustre-del-Rio 2014). The issue could thus be a lack of high-wage job opportunities (Moscarini and Postel-Vinay 2014). Such a development may be structural.[16]

REFERENCES

Abel, Jaison R., Richard Deitz, and Yaqin Su. 2014. "Are Recent College Graduates Finding Good Jobs?" Federal Reserve Bank of New York *Current Issues in Economics and Finance* 20 (1): 1–8.

Akerlof, George A., William T. Dickens, and George L. Perry. 1996. "The Macroeconomics of Low Inflation." *Brookings Papers on Economic Activity* 27 (1): 1–76.

Anger, Silke. 2011. "The Cyclicality of Effective Wages within Employer-Employee Matches in a Rigid Labor Market." *Labour Economics* 18 (December): 786–97.

Autor, David. 2010. "The Polarization of Job Opportunities in the U.S. Labor Market. Implications for Employment and Earnings." Available at http://economics.mit.edu/files/5554.

Barro, Robert J. 1980. "Discussion of Robert E. Hall's 'Labor Supply and Aggregate Fluctuations.'" *Carnegie-Rochester Conference Series on Public Policy* 12 (January): 35–8.

[16] See, for example, Autor (2010), Jaimovich and Siu (2012), and other works on job polarization.

Barrow, Lisa, and Jason Faberman. 2015. “Wage Growth, Inflation and the Labor Share.” Federal Reserve Bank of Chicago *Chicago Fed Letter* 349 (2015): 1–4.

Basu, Susanto, and Christopher L. House. 2015. “Challenges for New Keynesian Models with Sticky Wages.” Available at https://bfi.uchicago.edu/sites/default/files/research/Chicago_Handbook_Basu_House.pdf.

Bils, Mark J. 1985. “Real Wages over the Business Cycle: Evidence from Panel Data.” *Journal of Political Economy* 93 (August): 666–89.

Bils, Mark, Yongsung Chang, and Sun-Bin Kim. 2013. “How Sticky Wages In Existing Jobs Can Affect Hiring.” University of Rochester Working Paper.

Canon, Maria E., Marianna Kudlyak, Guannan Luo, and Marisa Reed. 2015. “Flows To and From Working Part Time for Economic Reasons and the Labor Market Aggregates in the Aftermath of the 2007–09 Recession.” Federal Reserve Bank of Richmond *Economic Quarterly* 100 (Second Quarter): 87–111.

Daly, Mary C., and Bart Hobijn. 2014. “Downward Nominal Wage Rigidities Bend the Phillips Curve.” Federal Reserve Bank of San Francisco Working Paper 2013-08.

Daly, Mary C., and Bart Hobijn. 2015. “Why Is Wage Growth So Slow?” Federal Reserve Bank of San Francisco *Economic Letter* 2015-01 (January 5).

Daly, Mary C., Bart Hobijn, and Brian Lucking. 2012. “Why has Wage Growth Stayed Strong?” Federal Reserve Bank of San Francisco *Economic Letter* 2012-10 (April 2).

Daly, Mary C., Bart Hobijn, and Theodore S. Wiles. 2012. “Dissecting Aggregate Real Wage Fluctuations: Individual Wage Growth and the Composition Effect.” Federal Reserve Bank of San Francisco Working Paper 2011-23.

Davis, Steven J., and John Haltiwanger. 2014. “Labor Market Fluidity and Economic Performance.” Working Paper 20479. Cambridge, Mass.: National Bureau of Economic Research. (September).

Elsby, Michael W. L., Donggyun Shin, and Gary Solon. 2014. “Wage Adjustment in the Great Recession and Other Downturns: Evidence from the United States and Great Britain.” Mimeo, University of Edinburg.

Foster, Lucia, John Haltiwanger, and C. J. Krizan. 2001. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." In *New Developments in Productivity Analysis*, edited by Edward Dean, Michael Harper, and Charles Hulten. Chicago: University of Chicago Press, 303–72.

Haefke, Christian, Marcus Sonntag, and Thijs van Rens. 2013. "Wage Rigidity and Job Creation." *Journal of Monetary Economics* 60 (November): 887–99.

Hobijn, Bart, and Leila Bengali. 2014. "The Wage Growth Gap for Recent College Grads." Federal Reserve Bank of San Francisco *Economic Letter* 2014-22 (July 21).

Jaimovich, Nir, and Henry E. Siu. 2012. "The Trend is the Cycle: Job Polarization and Jobless Recoveries." Working Paper 18334. Cambridge, Mass.: National Bureau of Economic Research. (August).

Kahn, Lisa B. 2010. "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy." *Labour Economics* 17 (April): 303–16.

Kudlyak, Marianna. 2009. "Are Wages Rigid Over the Business Cycle?" Federal Reserve Bank of Richmond *Economic Quarterly* 96 (Second Quarter): 179–99.

Kudlyak, Marianna. 2014. "The Cyclicality of the User Cost of Labor." *Journal of Monetary Economics* 68 (November): 53–67.

Lazear, Edward P., Kathryn L. Shaw, and Christopher Stanton. 2013. "Making Do With Less: Working Harder During Recessions." Working Paper 19328. Cambridge, Mass.: National Bureau of Economic Research. (August).

Martins, Pedro S., Gary Solon, and Jonathan P. Thomas. 2012. "Measuring What Employers Do about Entry Wages over the Business Cycle: A New Approach." *American Economic Journal: Macroeconomics* 4 (4): 36–55.

McLaughlin, Kenneth J., and Mark Bils. 2001. "Interindustry Mobility and the Cyclical Upgrading of Labor." *Journal of Labor Economics* 19 (January): 94–135.

Mester, Loretta J., and Elif Sen. 2013. "Has Job Quality Been 'Job One' in the Economic Recovery?" Federal Reserve Bank of Philadelphia *Research Rap*, 1–18. (August).

Moscarini, Giuseppe, and Fabien Postel-Vinay. 2014. "Did the Job Ladder Fail after the Great Recession?" Working Paper.

Mustre-del-Rio, Jose. 2014. "Following the Leaders: Wage Growth of Job Switchers." Federal Reserve Bank of Kansas City *Macro Bulletin* 2014 (December): 1–2.

Oi, Walter Y. 1962. "Labor as a Quasi-Fixed Factor." *Journal of Political Economy* 70 (December): 538–55.

Pissarides, Christopher A. 2009. "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?" *Econometrica* 77 (September): 1,339–69.

Solon, Gary, Robert Barsky, and Jonathan A. Parker. 1994. "Measuring the Cyclicality of Real Wages: How Important is Composition Bias." *The Quarterly Journal of Economics* 109 (February): 1–25.

Stiglitz, Joseph E. 1974. "Alternative Theories of Wage Determination and Unemployment in LDC's: The Labor Turnover Model." *The Quarterly Journal of Economics* 88 (May): 194–227.

Economic Quarterly—Volume 101, Number 3—Third Quarter 2015—Pages 245–274

Relative Price Changes and the Optimal Inflation Rate

Alexander L. Wolman

Relative prices of some goods or sectors have long-run trends: For example, the price of services relative to goods has been rising fairly steadily for decades. Other relative prices do not have long-run trends but sometimes fluctuate dramatically from one period to the next: For example, the relative price of the energy goods and services component of personal consumption expenditures fell by approximately 22 percent from January 2014 to January 2015. Although monetary policy changes can affect relative prices, in most cases trends in relative prices and dramatic monthly fluctuations in relative prices seem best viewed as exogenous factors with respect to monetary policy. How should monetary policy behave in the face of these fluctuations? More precisely, how is the optimal rate of inflation affected by exogenous relative price fluctuations? The answer, of course, depends on many factors, not least of which are the sources of monetary non-neutrality. We focus here on costly price adjustment as the principal source of non-neutrality. We use a two-good macroeconomic model with costly price adjustment to study the optimal constant inflation rate when there are trends in relative prices. For the purpose of studying short-run fluctuations in relative prices, we take an informal empirical approach, using one part of the model to construct hypothetical U.S. inflation rates that would have minimized the costs of price adjustment implied by the model. We also consider the potential conflict between minimizing the costs of price adjustment and meeting the central bank's announced inflation target.

The topic of relative price changes and optimal inflation has received much attention from researchers. In the New Keynesian litera-

The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.
E-mail: Alexander.Wolman@rich.frb.org.

ture, Aoki (2001) was an important early paper that made the point that goods with flexible prices should bear the burden of relative price adjustment. Aoki provides an argument for stabilizing core inflation related to the part of this paper that studies short-run relative price fluctuations. Bodenstein, Erceg, and Guerrieri (2008) is a related contribution. Neither Aoki nor Bodenstein, Erceg, and Guerrieri consider the effects of trend changes in relative prices, whereas Wolman (2008) provides an overview of this issue, and Wolman (2011) goes into more detail using models with staggered price setting and fixed costs of price adjustment. None of these previous papers construct historical time series for optimal inflation as a function of observable relative price changes, as we do in Section 5.

Both Aoki and Bodenstein use the Calvo price-adjustment framework, and Wolman (2011) also uses models with infrequent price adjustment. In contrast, this paper uses the Rotemberg model. An important factor in choosing Rotemberg instead of Calvo pricing is my desire to examine the implications of costly price adjustment, rather than infrequent price adjustment, for the optimal inflation rate. Thus, instead of calibrating the price-adjustment parameter to empirical work on the frequency of price adjustment, we calibrate it to Levy et al.'s (1997) estimates of the costs of price adjustment. Both the Calvo and Rotemberg approaches are reduced-forms, but the Rotemberg framework seems likely to be more appropriate as a stand-in for various *costs* of price change than the Calvo framework.[1] Additionally, the Rotemberg model is analytically quite simple, even when we consider trends in relative prices—nowhere in the paper do we need to rely on approximations in solving the model. That said, the qualitative results presented here are likely to hold up in the Calvo model or other staggered price-setting models.

The paper proceeds as follows. In Section 2 we describe the model in full detail, while in Section 3 we derive the conditions that characterize an equilibrium in which relative prices are diverging at a constant rate because of trend differences in productivity growth. In Section 4 we show how the optimal steady-state inflation rate varies with relative productivity growth and relative price-adjustment costs. In Section 5 we study optimal inflation on a period-by-period basis in a version of the model calibrated so that the two sectors represent energy and all other goods and services. For this exercise we use an ad-hoc static

[1] One example of a cost of price change for which the quadratic costs could be standing in is confusion created on the part of shoppers when nominal prices change. In the Calvo model, the costs of inflation are associated not with changes in prices but with variation in prices across goods at a point in time.

criterion that takes into account only current price-adjustment costs in each period. Section 6 concludes.

1. MODEL

In this section we will describe a complete macroeconomic model in which there are a large number of households, a large number of firms, and a monetary authority. The households are all identical; they consume the many goods produced by firms, which we will group into two categories called sector 1 and sector 2, and they supply labor that the firms use to produce those goods. Firms are monopolistically competitive, each one producing a unique good. Each firm has monopoly power for the good it produces, but it faces competition from the many other goods that are close substitutes according to households' preferences. Firms produce consumption goods using labor only, and productivity may be increasing steadily over time, at different rates for the two sectors. Firms set the dollar price for their goods, and there is a cost of changing price from one period to the next. The monetary authority chooses the rate of inflation, that is the change in the price index for the household's consumption basket. Unlike much of the quantitative dynamic stochastic general equilibrium literature, for the sake of simplicity the model does not incorporate any mechanisms to generate strategic complementarities among price-setting firms.

Households

The representative household is infinitely lived, discounts future utility at rate β, and has preferences over consumption (c_t) and labor supply (h_t) of two composite goods ($c_{k,t}$, $k = 1, 2$), and labor supply (h_t) given by

$$\sum_{t=0}^{\infty} \beta^t \left(\ln c_t - \chi h_t \right), \tag{1}$$

where the consumption basket c_t aggregates two sectors or categories ($c_{k,t}$, $k = 1, 2$),

$$c_t = c_{1,t}^{\nu} c_{2,t}^{1-\nu},$$

and each category is itself a composite of a continuum of differentiated products ($c_{k,t}(z)$):

$$c_{k,t} = \left(\int_0^1 c_{k,t}(z)^{\frac{\varepsilon-1}{\varepsilon}} dz \right)^{\frac{\varepsilon}{\varepsilon-1}}, \; k = 1, 2. \tag{2}$$

According to (2), consumers have a preference for variety—they would prefer to smooth their consumption of category k goods across all the goods in that category. There is a competitive labor market in which the real wage is w_t per unit of time. Households own the firms and receive a total nominal dividend payment from firms of Π_t.

The household's budget constraint is as follows, in dollar terms:

$$P_t c_t = P_t w_t h_t + \Pi_t,$$

or, in real terms:

$$c_t = w_t h_t + \frac{\Pi_t}{P_t}, \tag{3}$$

where P_t is the overall price level.

Optimality Conditions

From the utility function and the budget constraint, the household's intratemporal first-order conditions representing optimal choice of labor input and consumption are given by

$$\lambda_t w_t = \chi, \tag{4}$$

and

$$\lambda_t = 1/c_t, \tag{5}$$

where λ_t is the Lagrange multiplier on the budget constraint (3)—it can also be thought of as the marginal utility of an additional unit of consumption at time t. In (4), the income from an additional unit of time devoted to work is equated to the utility cost of that time. And in (5) the marginal utility of consumption is equated to the value of those resources in other uses.

Recall that overall consumption c_t is composed of consumption from two sectors, $c_{1,t}$ and $c_{2,t}$, and sectoral consumption is an aggregate of a continuum of differentiated goods. Although the consumer's problem does not involve subperiods, it can be useful to think of the consumer as first choosing overall consumption (c_t), then, given c_t, choosing the optimal split between $c_{1,t}$ and $c_{2,t}$, and finally, given those sectoral consumptions, choosing the optimal allocation across differentiated products within each sector ($c_{k,t}(z)$).

The optimality condition for the level of overall consumption is (5). The optimal choice of consumption from each sector minimizes the cost of consuming c_t given the sectoral price indices $P_{1,t}$ and $P_{2,t}$:

$$Min_{c_{1,t},c_{2,t}} \{P_{1,t} c_{1,t} + P_{2,t} c_{2,t}\} \tag{6}$$

$$s.t.\ c_t = c_{1,t}^{\nu} c_{2,t}^{1-\nu}. \tag{7}$$

The multiplier on the constraint in (6) is the nominal marginal cost of an additional unit of consumption—the price level. Therefore we denote the multiplier as P_t, and the first-order conditions are as follows:

$$\frac{c_{1,t}}{c_{2,t}} = \nu^{1/(1-\nu)} \left(\frac{P_{1,t}}{P_t}\right)^{1/(\nu-1)} \tag{8}$$

$$\frac{c_{1,t}}{c_{2,t}} = (1-\nu)^{-1/\nu} \left(\frac{P_{2,t}}{P_t}\right)^{1/\nu}. \tag{9}$$

Analogously, the optimal allocation of consumption within each sector minimizes the cost of consuming $c_{1,t}$ and $c_{2,t}$ given the prices of the differentiated products, and the multiplier on the constraint is the sectoral price index—the nominal marginal cost of purchasing one unit of the sectoral composite good:

$$Min_{c_{k,t}(z)} \left(\int_0^1 P_{k,t}(z)\, c_{k,t}(z)\, dz\right) \tag{10}$$

$$s.t.\ c_{k,t} = \left(\int_0^1 c_{k,t}(z)^{\frac{\varepsilon-1}{\varepsilon}}\, dz\right)^{\frac{\varepsilon}{\varepsilon-1}},\ k = 1, 2. \tag{11}$$

The first order conditions are as follows:

$$\frac{c_{k,t}(z)}{c_{k,t}} = \left(\frac{P_{k,t}(z)}{P_{k,t}}\right)^{-\varepsilon},\ k = 1, 2;\quad z \in (0, 1). \tag{12}$$

The Overall Price Level and Category Price Levels

Above, we derived the optimality conditions describing how households allocate consumption across sectors and across goods within each sector. Here we use those optimality conditions to show how prices within each sector aggregate to sectoral price indices and how the sectoral price indices aggregate to the overall price level.

From (11) and (12) we have

$$c_{k,t} = \left(\int_0^1 \left(\frac{P_{k,t}(z)}{P_{k,t}}\right)^{1-\varepsilon} c_{k,t}^{\frac{\varepsilon-1}{\varepsilon}}\, dz\right)^{\frac{\varepsilon}{\varepsilon-1}},\ k = 1, 2,$$

which can be manipulated to yield an expression for each sectoral price index as a function of individual goods prices within the sector:

$$P_{k,t} = \left(\int_0^1 (P_{k,t}(z))^{1-\varepsilon}\, dz\right)^{\frac{1}{1-\varepsilon}},\ k = 1, 2.$$

This expression for the category price indices plays an important role in staggered price-setting models, where there is equilibrium dispersion in the prices of like goods (varieties). Here however, all prices within a category will be identical in the equilibria we consider.[2]

Across sectors, if productivity is not identical then there will be price differentials. From (8) and (9) we have

$$(1-\nu)^{-1/\nu}\left(\frac{P_{2,t}}{P_t}\right)^{1/\nu}=\nu^{1/(1-\nu)}\left(\frac{P_{1,t}}{P_t}\right)^{1/(\nu-1)}, \qquad (13)$$

which can be manipulated to yield an expression for the price level as a function of the sectoral price indices:

$$P_t=\left(\frac{P_{1,t}}{\nu}\right)^{\nu}\left(\frac{P_{2,t}}{1-\nu}\right)^{1-\nu}. \qquad (14)$$

Firms

Each firm z in sector k produces consumption goods with a technology that is linear in labor:

$$y_{k,t}(z)=a_{k,t}h_{k,t}(z).$$

In equilibrium the firm's output will be identical to households' consumption of that good ($y_{k,t}(z)=c_{k,t}(z)$), but for the purposes of this section we will maintain a distinction between output and consumption. Productivity ($a_{k,t}$) may vary across time and across sectors, but it is the same across firms within a sector. Marginal cost for each firm in sector k ($\psi_{k,t}$) is given by the ratio of the wage to the marginal product of labor:

$$\psi_{k,t}=w_t/a_{k,t}.$$

Firms face a cost $\left(\xi_{k,t}(z)\right)$ in terms of labor of changing the nominal price of the good they produce (z). The cost is proportional to output and is affected by the same productivity shifter as regular goods production:

$$\xi_{k,t}(z)=\frac{\theta_k}{2}\left(\frac{y_{k,t}(z)}{a_{k,t}}\right)\left(\frac{P_{k,t}(z)}{P_{k,t-1}(z)}-1\right)^2. \qquad (15)$$

An individual firm chooses its price each period to maximize the expected present value of profits, where profits in any single period

[2] In principle one could imagine equilibria in which initial dispersion in prices was sustained over time, but these equilibria have not been the focus of attention in Rotemberg models.

are given by revenue minus costs of production minus costs of price adjustment, plus a per-unit subsidy from the government.[3]

From (12), the demand curve facing firm z in sector k is $y_{k,t}(z) = (P_{k,t}(z)/P_{k,t})^{-\varepsilon} y_{k,t}$, so the profit maximization problem for such a firm is

$$\max_{P_{k,t+j}(z)} \sum_{j=0}^{\infty} \beta^j \left(\frac{\lambda_{t+j}}{\lambda_t}\right) \left[(1+\tau) \frac{P_{k,t+j}(z)}{P_{t+j}} \left(\frac{P_{k,t+j}(z)}{P_{k,t+j}}\right)^{-\varepsilon} y_{k,t+j} \right.$$
$$- \frac{w_{t+j}}{a_{k,t+j}} \left(\frac{P_{k,t+j}(z)}{P_{k,t+j}}\right)^{-\varepsilon} y_{k,t+j}$$
$$\left. - \frac{\theta_k}{2} w_{t+j} \left(\left(\frac{P_{k,t+j}(z)}{P_{k,t+j}}\right)^{-\varepsilon} \frac{y_{k,t+j}}{a_{k,t+j}} \right) \left(\frac{P_{k,t+j}(z)}{P_{k,t+j-1}(z)} - 1 \right)^2 \right].$$

The first term in the square brackets is the real revenue a firm earns charging price $P_{t+j}(z)$ in period $t+j$; it sells $(P_{k,t+j}(z)/P_{k,t+j})^{-\varepsilon} y_{k,t+j}$ units of goods for relative price $P_{t+j}(z)/P_{t+j}$, and it receives a subsidy of τ per unit sold. The second term represents the cost of producing those goods, and the third term represents the price-adjustment cost.

Note that the price chosen in any period shows up only in two periods of the infinite sum. Thus, the part of the objective function relevant for the choice of a price in period t is

$$(1+\tau) \frac{P_{k,t}(z)}{P_t} \left(\frac{P_{k,t}(z)}{P_{k,t}}\right)^{-\varepsilon} y_{k,t} - \frac{w_t}{a_{k,t}} \left(\frac{P_{k,t}(z)}{P_{k,t}}\right)^{-\varepsilon} y_{k,t}$$
$$- \frac{\theta_k}{2} w_t \left(\frac{(P_{k,t}(z)/P_{k,t})^{-\varepsilon} y_{k,t}}{a_{k,t}} \right) \left(\frac{P_{k,t}(z)}{P_{k,t-1}(z)} - 1 \right)^2$$
$$- \frac{\lambda_{t+1}}{\lambda_t} \frac{\beta\theta_k}{2} w_{t+1} \left(\frac{(P_{k,t+1}(z)/P_{k,t+1})^{-\varepsilon} y_{k,t+1}}{a_{k,t+1}} \right) \left(\frac{P_{k,t+1}(z)}{P_{k,t}(z)} - 1 \right)^2.$$

[3] The purpose of including the subsidy is to allow us to focus attention mainly on the direct costs of price adjustment, rather than secondary effects through changes in markups.

The first-order condition is:

$$
\begin{aligned}
0 \quad = \quad & \frac{(1+\tau)(1-\varepsilon)}{P_t}\left(\frac{P_{k,t}(z)}{P_{k,t}}\right)^{-\varepsilon} y_{k,t}+\frac{\varepsilon w_t}{a_{k,t}P_{k,t}}\left(\frac{P_{k,t}(z)}{P_{k,t}}\right)^{-\varepsilon-1} y_{k,t} \\
& -\theta_k w_t\left(\frac{\left(P_{k,t}(z)/P_{k,t}\right)^{-\varepsilon} y_{k,t}}{a_{k,t}}\right)\left(\frac{P_{k,t}(z)}{P_{k,t-1}(z)}-1\right)\frac{1}{P_{k,t-1}(z)} \\
& +\frac{\varepsilon\theta_k}{2} w_t\left(\frac{1}{a_{k,t}}\right)\frac{1}{P_{k,t}}\left(\frac{P_{k,t}(z)}{P_{k,t}}\right)^{-\varepsilon-1} y_{k,t}\left(\frac{P_{k,t}(z)}{P_{k,t-1}(z)}-1\right)^2 \\
& +\frac{\lambda_{t+1}}{\lambda_t}\beta\theta_k w_{t+1}\frac{y_{k,t+1}}{a_{k,t+1}}\left(\frac{P_{k,t+j}(z)}{P_{k,t+j}}\right)^{-\varepsilon}\left(\frac{P_{k,t+1}(z)}{P_{k,t}(z)}-1\right)\frac{P_{k,t+1}(z)}{\left(P_{k,t}(z)\right)^2}.
\end{aligned}
$$

Absent price-adjustment costs (if θ_k were zero) the firm would choose price in order to balance the marginal revenue and marginal cost of further price adjustment, yielding a constant markup over marginal cost: $P_{k,t}(z) = (\varepsilon/((\varepsilon-1)(1+\tau)))(w_t/a_{k,t})P_t$. With price-adjustment costs, the same principle applies, but marginal cost is now more complicated, because price adjustment also affects (marginal) price-adjustment costs in the current and subsequent periods. There is no closed form solution for the optimal price, although we will see below that in a symmetric equilibrium the expression simplifies dramatically.

Monetary Authority

In reality, the monetary authority controls instrument variables, like the size of its balance sheet or some administered interest rates. Inflation is then an equilibrium outcome. In the model we assume that the monetary authority directly controls the inflation rate: $\pi_t = P_t/P_{t-1}$. How central banks determine the inflation rate, and how precisely they can control the inflation rate, are interesting and important questions, but they are beyond the scope of this paper.

Equilibrium Conditions

We consider a symmetric equilibrium where all firms in the same sector charge the same price and hence produce the same quantity. Thus we have

$$\left.\begin{array}{l} P_{k,t}(z) = P_{k,t} \\ y_{k,t}(z) = y_{k,t} \end{array}\right\}, \quad k = 1,2. \tag{16}$$

And because goods are produced for consumption only, in a symmetric equilibrium with all firms in sector k producing the same quantity, the goods market-clearing condition for each sector is

$$y_{k,t} = c_{k,t}. \tag{17}$$

The labor market-clearing condition states that labor supplied by households (h_t) must be equal to the sum of labor used in production and labor used in price adjustment. Because of symmetry, we can just refer to the sectoral aggregates instead of individual firm inputs:

$$h_t = h_{1,t} + h_{2,t} + \xi_{1,t} + \xi_{2,t}.$$

Using the price-adjustment cost function (15) and the technology for goods production, and imposing (16), we have

$$h_t = \frac{y_{1,t}}{a_{1,t}} \cdot \left(1 + \frac{\theta_1}{2}\left(\pi_{1,t} - 1\right)^2\right) + \frac{y_{2,t}}{a_{2,t}} \cdot \left(1 + \frac{\theta_2}{2}\left(\pi_{2,t} - 1\right)^2\right), \quad (18)$$

where $\pi_{k,t}$ denotes the gross rate of sectoral price change ($\pi_{k,t} = P_{k,t}/P_{k,t-1}$).

Imposing symmetry on the firms' first-order condition yields the following expression, which we will refer to as the sectoral Phillips curve:

$$\begin{aligned} 0 \;=\; & (1+\tau)(1-\varepsilon)\, p_{k,t} y_{k,t} + \varepsilon w_t \frac{y_{k,t}}{a_{k,t}} \\ & -\theta_k w_t \left(\frac{y_{k,t}}{a_{k,t}}\right)(\pi_{k,t} - 1)\, \pi_{k,t} \\ & + \frac{\varepsilon \theta_k}{2} w_t \left(\frac{y_{k,t}}{a_{k,t}}\right)(\pi_{k,t} - 1)^2 \\ & + \frac{\lambda_{t+1}}{\lambda_t} \beta \theta_k w_{t+1} \left(\frac{y_{k,t+1}}{a_{k,t+1}}\right)(\pi_{k,t+1} - 1)(\pi_{k,t+1}), \end{aligned}$$

where we use $p_{k,t}$ to denote the relative price of the sector k good: $p_{k,t} \equiv P_{k,t}/P_t$. The other equilibrium conditions are the household's first-order conditions (4), (5), (8), and (9) and the price level equation (14), which can also be written as a restriction on relative prices:

$$1 = \left(\frac{p_{1,t}}{\nu}\right)^{\nu} \left(\frac{p_{2,t}}{1-\nu}\right)^{1-\nu}; \quad (19)$$

an increase in the relative price of one good implies a decrease in the relative price of the other good.

2. A STEADY STATE WITH TRENDING RELATIVE PRICES

Thus far, we have not specified the processes for sectoral productivity $a_{k,t}$, which, along with the inflation rate, are the model's driving variables—that is, the exogenous variables that determine the nature of equilibrium outcomes. In this section we impose constant growth rates for productivity and derive the equilibrium values of the model's endogenous variables as functions of parameters and the rates of productivity growth and inflation. Then we compute some examples for particular parameter values. We will only consider an equilibrium in which all growth rates are constant and where labor input is constant in both sectors.[4]

The System of Equations

Assume that productivity grows at constant gross rate μ_k in sector k, and make the normalizing assumption that in period zero the level of productivity in both sectors was 1. Then we have

$$a_{k,t} = \mu_k^t.$$

If we also assume that the inflation rate is constant, equal to π, then we can write the steady-state Phillips curve equations as follows (at this point we are leaving t subscripts on all variables, because it is not yet obvious which variables will be constant in a steady state with growth and how to normalize all variables in that steady state):

$$\begin{aligned} 0 \quad = \quad & (1+\tau)(1-\varepsilon) p_{k,t} y_{k,t} + \varepsilon \frac{w_t}{\mu_k^t} y_{k,t} \\ & -\theta_k w_t \left(\frac{y_{k,t}}{\mu_k^t}\right)\left(\frac{p_{k,t}}{p_{k,t-1}}\pi - 1\right)\frac{p_{k,t}}{p_{k,t-1}}\pi \\ & +\frac{\varepsilon\theta_k}{2} w_t \left(\frac{1}{\mu_k^t}\right) y_{k,t} \left(\frac{p_{k,t}}{p_{k,t-1}}\pi - 1\right)^2 \\ & +\frac{\lambda_{t+1}}{\lambda_t}\beta\theta_k w_{t+1}\left(\frac{y_{k,t+1}}{\mu_k^{t+1}}\right)\left(\frac{p_{k,t+1}}{p_{k,t}}\pi - 1\right)\left(\frac{p_{k,t+1}}{p_{k,t}}\pi\right), \quad k = 1, 2. \end{aligned} \tag{20}$$

From the household's first-order conditions, (4), (5), (8), and (9), the definition of the consumption index (7), and the goods market clearing condition (17), we can express output in the two sectors as a function of the wage and relative prices:

$$y_{1,t} = \frac{\nu}{\chi}\frac{w_t}{p_{1,t}}, \tag{21}$$

$$y_{2,t} = \frac{1-\nu}{\chi}\frac{w_t}{p_{2,t}}. \tag{22}$$

[4] We have not proved that the equilbrium we study is unique, though we suspect that it is.

Finally, we use these last two equations to eliminate sectoral outputs in the Phillps curve. After simplifying, we have

$$
\begin{aligned}
0 \quad = \quad & (1+\tau)(1-\varepsilon)+\varepsilon \frac{w_t}{\mu_1^t p_{1,t}} \\
& -\theta_1\left(\frac{w_t}{\mu_1^t}\right)\left(\frac{p_{1,t}}{p_{1,t-1}}\pi-1\right)\frac{\pi}{p_{1,t-1}} \\
& +\frac{\varepsilon\theta_1}{2}\frac{w_t}{\mu_1^t p_{1,t}}\left(\frac{p_{1,t}}{p_{1,t-1}}\pi-1\right)^2 \\
& +\beta\theta_1\left(\frac{w_{t+1}}{\mu_1^{t+1}p_{1,t+1}}\right)\left(\frac{p_{1,t+1}}{p_{1,t}}\pi-1\right)\left(\frac{p_{1,t+1}}{p_{1,t}}\pi\right).
\end{aligned}
$$

Analogously, for sector 2 the Phillips curve is as follows:

$$
\begin{aligned}
0 \quad = \quad & (1+\tau)(1-\varepsilon)+\varepsilon \frac{w_t}{\mu_2^t p_{2,t}} \\
& -\theta_2\left(\frac{w_t}{\mu_2^t}\right)\left(\frac{p_{2,t}}{p_{2,t-1}}\pi-1\right)\frac{\pi}{p_{2,t-1}} \\
& +\frac{\varepsilon\theta_2}{2}\frac{w_t}{\mu_2^t p_{2,t}}\left(\frac{p_{2,t}}{p_{2,t-1}}\pi-1\right)^2 \\
& +\beta\theta_2\left(\frac{w_{t+1}}{\mu_2^{t+1}p_{2,t+1}}\right)\left(\frac{p_{2,t+1}}{p_{2,t}}\pi-1\right)\left(\frac{p_{2,t+1}}{p_{2,t}}\pi\right).
\end{aligned}
$$

Together with the restriction on relative prices implied by the price index (19), these two Phillips curve equations represent a system of three equations in the three endogenous variables $w_t, p_{1,t}$, and $p_{2,t}$. In order to make further progress however, we need to use the constant productivity growth assumption to find the growth rates of wages and relative prices.

Productivity Growth Determines Growth Rates of Wages and Relative Prices

From the household's optimality conditions, (4) and (5), we know that $w_t = \chi c_t$, so the real wage must grow at the same rate as consumption. What do we know about the growth rate of c? Let γ_{c_k} be the gross growth rate of c_k. Then, if we denote the growth rate of consumption as γ_c, we have $\gamma_c = \gamma_{c_1}^{\nu}\gamma_{c_2}^{1-\nu}$:

$$
\underbrace{c_t}_{\gamma_c} = \left(\underbrace{c_{1,t}}_{\gamma_{c_1}}\right)^{\nu}\left(\underbrace{c_{2,t}}_{\gamma_{c_2}}\right)^{1-\nu},
$$

Now consider the growth rates of c_1 and c_2. With labor input constant in the two sectors, output grows at the rate of productivity, so we know the growth rates of c_1, c_2, c, and w :

$$\begin{aligned}\gamma_{c_1} &= \mu_1 \\ \gamma_{c_2} &= \mu_2 \\ \gamma_c &= \mu_1^{\nu}\mu_2^{1-\nu} \\ \gamma_w &= \mu_1^{\nu}\mu_2^{1-\nu}.\end{aligned}$$

To determine the growth rates of relative prices we use this same reasoning on (21) and (22):

$$\begin{aligned}\nu \underbrace{w_t}_{\mu_1^{\nu}\mu_2^{1-\nu}} &= \chi \underbrace{p_{1,t}}_{\gamma_{p_1}} \underbrace{c_{1,t}}_{\mu_1} \\ (1-\nu) \underbrace{w_t}_{\mu_1^{\nu}\mu_2^{1-\nu}} &= \chi \underbrace{p_{2,t}}_{\gamma_{p_2}} \underbrace{c_{2,t}}_{\mu_2}.\end{aligned}$$

That is

$$\gamma_{p_1} = (\mu_1/\mu_2)^{\nu-1} \tag{23}$$
$$\gamma_{p_2} = (\mu_1/\mu_2)^{\nu}. \tag{24}$$

Note that the growth rates of relative prices must cancel out when aggregated according to the price index (19), and the growth rates we just derived indeed have this property.

The Time Paths of Wages and Relative Prices

We have now determined the growth rates of wages and relative prices, but we need to use the price index and Phillips curve equations to determine the levels of those variables. To that end, we plug into the Phillips curves and the price index expressions for wages and relative prices that use the trends we just derived together with unknown initial levels:

$$\begin{aligned}w_t &= w_0 \cdot \left(\mu_1^{\nu}\mu_2^{1-\nu}\right)^t \\ p_{1,t} &= p_{1,0} \cdot (\mu_1/\mu_2)^{(\nu-1)t} \\ p_{2,t} &= p_{2,0} \cdot (\mu_1/\mu_2)^{\nu t}.\end{aligned}$$

The equations will then be used to determine the initial levels $w_0, p_{1,0}$, and $p_{2,0}$. Once we have simplified the Phillips curves and collected terms, we have

$$\frac{p_{1,0}}{w_0} = \frac{\varepsilon}{(\varepsilon-1)(1+\tau)} \times \tag{25}$$

$$\left\{ 1 - \theta_1 \left(\left(\frac{\mu_1}{\mu_2} \right)^{\nu-1} \pi - 1 \right) \times \right.$$

$$\left. \left[\left(\frac{1-\beta}{\varepsilon} \right) \left(\frac{\mu_1}{\mu_2} \right)^{\nu-1} \pi - \frac{1}{2} \left(\left(\frac{\mu_1}{\mu_2} \right)^{\nu-1} \pi - 1 \right) \right] \right\} \tag{26}$$

and analogously for good 2,

$$\frac{p_{2,0}}{w_0} = \frac{\varepsilon}{(\varepsilon-1)(1+\tau)} \times \tag{27}$$

$$\left\{ 1 - \theta_2 \left(\left(\frac{\mu_1}{\mu_2} \right)^{-\nu} \pi - 1 \right) \times \right.$$

$$\left. \left[\left(\frac{1-\beta}{\varepsilon} \right) \left(\frac{\mu_1}{\mu_2} \right)^{-\nu} \pi - \frac{1}{2} \left(\left(\frac{\mu_1}{\mu_2} \right)^{-\nu} \pi - 1 \right) \right] \right\}. \tag{28}$$

And the price index equation is

$$1 = \left(\frac{p_{1,0}}{\nu} \right)^{\nu} \left(\frac{p_{2,0}}{1-\nu} \right)^{1-\nu}. \tag{29}$$

So we have three equations in $p_{1,0}, p_{2,0}$, and w_0. For any values of the parameters $(\varepsilon, \nu, \beta, \mu_j, \theta_j,)$ and the inflation rate (π) it is straightforward to calculate the values of all the endogenous variables. First, $p_{1,0}, p_{2,0}$, and w come from solving the last three equations. The other important variables, c and h, can easily be calculated from (21), (22), and (18). Detailed derivations are contained in the Appendix.

3. OPTIMAL STEADY-STATE INFLATION

In this section we explore the model's implications for the optimal steady-state inflation rate, focusing on how differential productivity growth rates and differential price-adjustment costs in the two sectors affect the optimal inflation rate. Without that sectoral heterogeneity, the question would be relatively simple, although there is one subtlety that we will quickly dispense with. In a one-sector version of the model, or equivalently a model without heterogeneity in productivity growth, it would be possible to eliminate all direct, steady-state costs of price adjustment by maintaining a zero inflation rate: Prices would never need to change. However, in the absence of an appropriate subsidy τ, zero inflation would not be the welfare-maximizing steady-state inflation rate because a small amount of inflation would reduce the average markup of price over marginal cost toward its efficient level of zero. For our purposes, the latter effect is a nuisance—it detracts attention from the focus of the paper. So in what follows we will set the subsidy τ so that it exactly offsets the markup in a steady state with zero inflation and identical sectors: $\tau = 1/(\varepsilon - 1)$.

The model is simple enough that we can use the equations above to derive the level of steady-state welfare as an explicit function of inflation. However, the expressions become somewhat complicated, so we simply reiterate here that inflation enters the welfare calculation through two channels. First, there is the direct effect on price-adjustment costs (see [18]). Second, there is an effect on relative prices and the real wage, that can be seen in (25)–(29). This is the markup effect mentioned in the previous paragraph. For any values of the parameters and for a range of inflation rates, it would be possible to choose the subsidy to eliminate the markup in one sector, and the markup could be eliminated in both sectors with sector-specific subsidies. For simplicity we fix the subsidy at the level that eliminates the markup in a zero-inflation steady state with common productivity growth across sectors. Later on we will abstract entirely from the markup effects of inflation. However, here our calculations of optimal steady-state welfare will incorporate both the price-adjustment cost and markup effects.

To help interpret the figures below, recall that the overall price level is related to the sectoral price indices as follows,

$$P_t = \left(\frac{P_{1,t}}{\nu}\right)^{\nu} \left(\frac{P_{2,t}}{1-\nu}\right)^{1-\nu}. \tag{30}$$

The inflation rate, then, can be written as a function of the rates of price change in the two sectors:

$$\pi_t = \frac{P_t}{P_{t-1}} = \pi_{1,t}^{\nu} \pi_{2,t}^{1-\nu}. \tag{31}$$

Using the log approximation, for moderate rates of price change the net inflation rate is approximately equal to the share weighted average of the two sectoral rates of price change:

$$\pi_t \approx \nu \pi_{1,t} + (1-\nu)\,\pi_{1,t}. \tag{32}$$

Parameter Values

To study the effect of relative productivity growth and relative price-adjustment costs on the optimal inflation rate, we set the parameters β and ε at standard values in the literature, 0.99 and 10, respectively. Interpreting a model time period as one quarter, $\beta = 0.99$ implies an annual real interest rate of about 4 percent. With $\varepsilon = 10$, the average

Figure 1 Optimal Inflation with $v = 0.5$

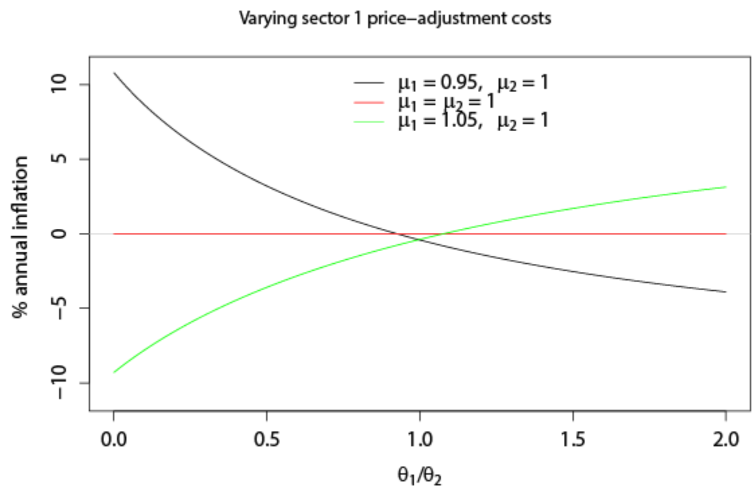

markup is around 11 percent absent the subsidy. As a benchmark we set the rates of productivity growth in each sector to 1 and the price-adjustment cost parameter (θ_j) in each sector to 140. With $\theta_1 = \theta_2 = 140$, the steady-state price-adjustment costs associated with a 1 percent quarterly price change amount to 0.7 percent of a firm's revenue. This is the average number estimated by Levy et al. (1997) in their study of supermarkets.

Relative Price Trends and The Optimal Inflation Rate

Figures 1 and 2 plot the optimal steady-state inflation rate as a function of relative price-adjustment costs (Figure 1) and relative productivity growth (Figure 2) for the case where consumers' expenditure shares are equal for the two sectors; that is, $\nu = 0.5$. In Figure 1, the black line represents a case with relatively low productivity growth in sector 1 ($\mu_1/\mu_2 = 0.95$), the red line is constant and equal productivity in both sectors ($\mu_1 = \mu_2 = 1$), and the green line is relatively high productivity growth in sector 1 ($\mu_1/\mu_2 = 1.05$). Since it is simplest, focus first on the red line. When productivity is constant in both sectors, relative prices are also constant, and when relative prices are constant there is no need for any nominal prices to change. Thus, since zero inflation

Figure 2 Optimal Inflation with $v = 0.5$

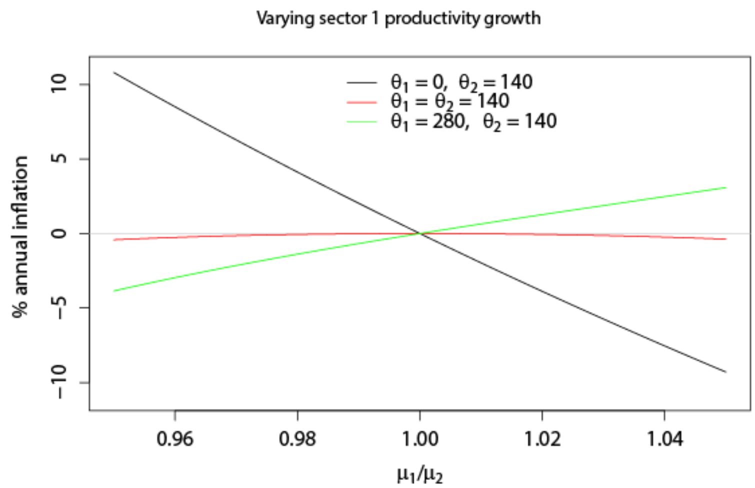

eliminates the costs of price change, the optimal inflation rate is zero because there are price-adjustment costs in sector 2; this holds regardless of the level of price-adjustment costs in sector 1. When productivity growth does vary across sectors, relative prices must change in steady state—specifically, the sector with low productivity growth should see its relative price rise. In order for relative prices to change, at least one nominal price needs to change. In the case when productivity growth is lower in sector 1 (black line), the optimal inflation rate is decreasing in sector 1's price-adjustment costs: When sector 1's adjustment costs are low (θ_1/θ_2 is low), it is optimal for the required increase in sector 1 relative prices to occur through an increase in sector 1 nominal prices because those price increases are not costly. In contrast, when θ_1/θ_2 is high, the increase in sector 1 relative prices optimally occurs mainly through a decrease in sector 2 nominal prices, because those price decreases are low cost. The explanation for the case of high sector 1 productivity growth (green line) essentially involves reversing the signs relative to the black line.

Figure 2 displays similar relationships, except that relative productivity growth varies continuously on the horizontal axis, and the three lines represent different levels of relative price-adjustment costs. The black and green lines represent low and high levels of sector 1 price-adjustment costs (θ_1), and the red line is the case where $\theta_1 = \theta_2$.

Figure 3 Optimal Inflation with $v = 0.05$

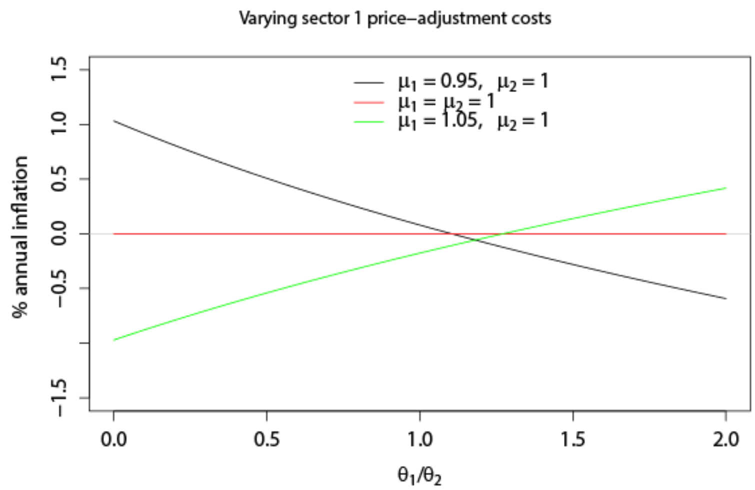

A curious feature of Figures 1 and 2 is that when price-adjustment costs are equal in the two sectors, a small amount of deflation is optimal. In Figure 1, the point where the black and green lines intersect exhibits deflation, and in Figure 2 the red line exhibits deflation everywhere except when productivity growth is equal in the two sectors. To understand this result, consider the case where there is positive productivity growth in sector 1 and zero productivity growth in sector 2 (i.e., $\mu_1 > 1, \mu_2 = 1$). Given the quadratic price-adjustment costs, zero overall inflation might seem optimal: Prices would be falling somewhat in sector 1 and rising somewhat in sector 2. The problem with zero inflation is that price-adjustment costs would actually be larger in sector 2 (with rising prices) than in sector 1. Price-adjustment costs can be reduced by shifting some of the price-adjustment burden toward sector 1, and this requires slight deflation. Formally, with $\nu = 0.5$, from (31) zero inflation implies $\pi_1 = 1/\pi_2$. Price-adjustment costs are then proportional to $(\pi_1 - 1)^2$ in sector 1 and proportional to $(1/\pi_1 - 1)^2$ in sector 2. With $\pi_1 < 1$, adjustment costs would be higher in sector 2.

Figures 3 and 4 are analogues to Figures 1 and 2 for the case where the expenditure share on sector 1 is small, just 5 percent. Note that the inflation magnitudes in these figures are much smaller. With one sector very small, nominal price changes in that sector pass through

Figure 4 Optimal Inflation with v = 0.05

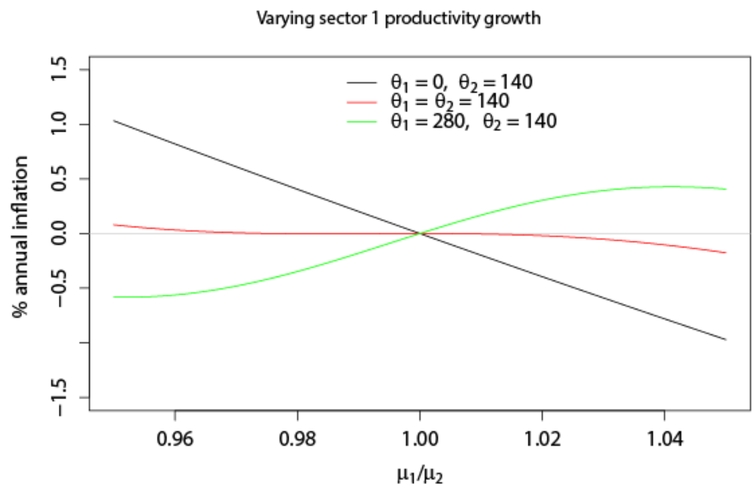

less to overall inflation (see [32]), so it is possible to achieve the desired relative price changes with less overall inflation (or deflation).

The figures above show how optimal inflation in a two-sector model varies across a wide range for relative productivity growth, relative price-adjustment costs, and relative size of the two sectors. Although the model is highly stylized, one could interpret the figures as indicating that the optimal inflation rate in the United States is actually negative. This follows from the fact that goods prices seem to be more flexible than services prices (Bils and Klenow 2004) and the price of goods relative to services has trended down over time. Optimality of deflation is not an uncommon result, but that result is typically associated with the shoe-leather costs of inflation (Friedman 1969).

4. SHORT-RUN RELATIVE PRICE CHANGES AND OPTIMAL INFLATION

The preceding analysis of optimal steady-state inflation is relevant for thinking about situations where there are long-run trends in relative prices across sectors. However, there are even larger short-run fluctuations in relative prices, and these can occur for goods or sectors that do not experience trend relative price changes. The obvious example is energy prices, which have fluctuated dramatically in the United States recently. In the presence of these large relative price movements, what is the optimal behavior of inflation period-by-period? The model laid out above can provide an answer to that question for any parameter values and any fluctuations in sectoral productivity growth—which are the natural drivers of relative prices in the model. In this section we will study an interesting special case: Sector 1 will be interpreted as *energy goods and services*, and sector 2 as all other personal consumption expenditure goods and services. We will assume that price-adjustment costs are zero for energy goods and services, whereas for all other goods those costs are determined by (15) with $\theta_2 = 140$, as above.[5] We will further assume that there are sector-specific and, in the case of sector 2, time-varying subsidies that make price always equal to marginal cost. Thus, the inflation rate affects only the magnitude of price-adjustment costs, given the exogenous fluctuations in relative prices.[6]

In the special case just described, optimal policy is trivial. The relative price of goods across the two sectors must change over time because of changes in relative productivity. With zero price-adjustment costs in sector 1, it is optimal for the entire burden of adjustment to be born by sector 1, and this can be accomplished in every period without any distortions. That is, the markup will be zero, and zero adjustment costs will be incurred. Such a policy is related to stabilizing core inflation, except that here core inflation refers to ex-energy inflation instead of ex-food and energy. In Section 5 we derive the counterfactual inflation rate for the United States since 1959 that would have minimized price-adjustment costs under these assumptions. We then incorporate the fact that since 2012 the Federal Reserve has had a 2 percent target for inflation. We calculate modified counterfactual inflation rates that balance a desire to meet the inflation target with a desire to minimize price-adjustment costs.

Time Path of Inflation Chosen to Minimize Price-Adjustment Costs

With zero price-adjustment costs in sector 1, the inflation rate that minimizes price-adjustment costs (call it π_t^a) is the inflation rate that accommodates all relative price changes with nominal price changes only for the sector 1 good. That is,

$$\frac{\pi_{1,t}^a}{\pi_{2,t}^a} = \pi_{1,t}^a = \frac{p_{1,t}/p_{1,t-1}}{p_{2,t}/p_{2,t-1}}$$

[5] This approach follows Aoki (2001), who used a model with Calvo price setting.

[6] Under the stated assumptions, the ratio of relative prices in the two sectors is given by the inverse of the ratio of productivities: $p_{2,t}/p_{1,t} = a_{1,t}/a_{2,t}$.

Figure 5 Inflation Rate to Minimize Price-Adjustment Costs

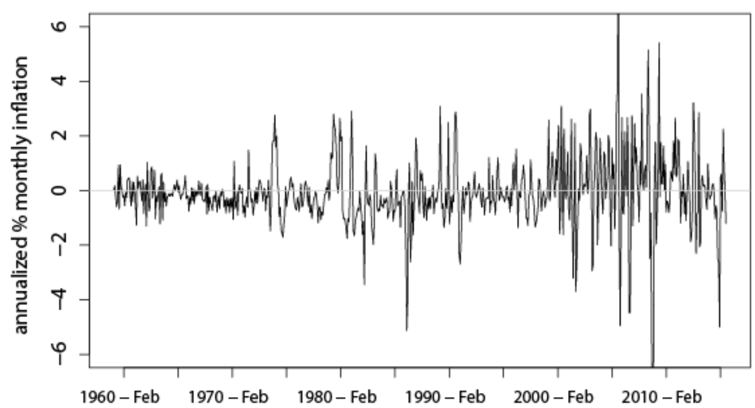

and,

$$\pi_t^a = \left(\pi_{1,t}^a\right)^{\nu},$$

which together imply

$$\pi_t^a = \left(\frac{p_{1,t}/p_{1,t-1}}{p_{2,t}/p_{2,t-1}}\right)^{\nu}. \tag{33}$$

With relative price changes viewed as exogenous, and thus equal to their observed values, we can use (33) to construct a time series for π_t^a. In constructing π_t^a we deviate from the assumptions above by allowing ν to vary each period, setting it equal to that period's observed energy expenditure share.

Figure 5 displays the monthly time series for π_t^a, the inflation rate that would have minimized price-adjustment costs under the stated assumptions. Relative prices for the energy sector are extremely volatile, and that volatility is optimally allocated entirely to nominal energy prices. Nominal energy price volatility then passes directly through to inflation. However, because the energy expenditure share is relatively low, the resulting fluctuations in inflation are much smaller than the fluctuations in energy prices. Note that since 2000 the volatility of π_t^a has increased markedly due to an increase in the volatility of the relative price of energy.

Figure 6 Inflation Rate to Minimize Price-Adjustment Costs, 12-month

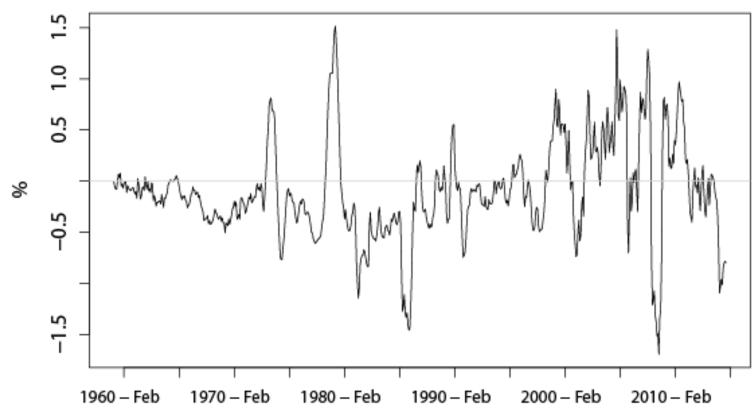

Figure 6 also plots π_t^a, but it displays that inflation rate over the preceding 12 months, instead of the monthly values. According to our simple model, for the most recently available data, adjustment costs would have been minimized with deflation of around 1.25 percent.

While Figures 5 and 6 emphasize the time series behavior of π_t^a, it is also interesting to look at the entire distribution of "optimal" inflation outcomes. Expressed as deviations from their mean, this distribution may be relevant in thinking about an appropriate band within which to target inflation. Figure 7 displays the sample distribution function of 12-month π_t^a, in deviations from its mean.

Approximately 90 percent of the time, the inflation rate that minimizes price adjustment costs lies between −1 percent and +1 percent, relative to its mean. This finding lends some tentative support to the notion that a reasonable inflation targeting band would be +/− 1 percentage point: If the costs of inflation are given by quadratic costs of price adjustment as calibrated here, then historical fluctuations in relative energy prices imply that it is optimal to keep inflation within a +/− 1 percentage point band 90 percent of the time. From 1959 to the present, inflation has been substantially more volatile than this hypothetical distribution, with 90 percent of the observations lying between −2.2 percent and 5.5 percent relative to the mean. Of course, there

Figure 7 Sample Distribution of Optimal Inflation Deviations From Mean

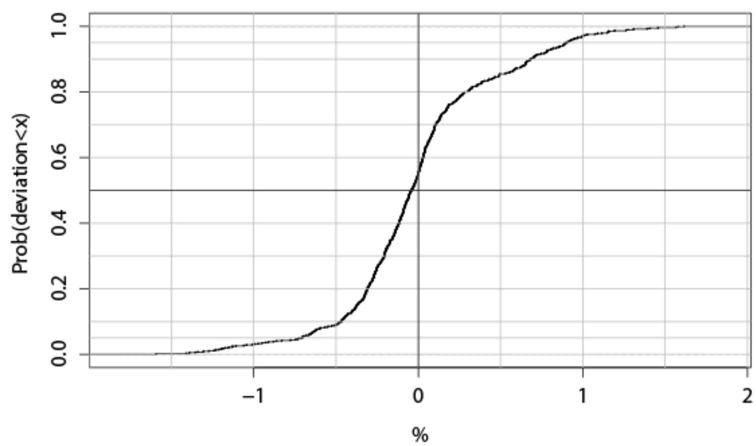

were large swings in trend inflation over that period, so it makes more sense to compare the hypothetical distribution to actual inflation in the inflation targeting period, which officially starts in January 2012. Over that period the volatility of 12-month inflation has actually been close to the volatility displayed in Figure 7: the 5th and 95th percentiles over this period are approximately −1.2 percent and 1.1 percent relative to the sample mean. Note however that these percentiles are based on a very short sample.

Inflation-Targeting Era: Trading Off Inflation Target Misses and Price-Adjustment Costs

The fact that the Federal Reserve's inflation target is 2 pecent instead of zero motivates an alternative notion of optimal inflation to the one used above. That alternative weights equally the inflation rate that minimizes adjustment costs and the 2 percent inflation target. The inflation rate that minimizes adjustment costs is close to zero on average, because there is no appreciable long-run trend in relative energy prices. Therefore, on average this alternative measure of optimal inflation will lie between 2 percent and zero. Figures 8 and 9 plot the implied inflation rate (measured in monthly and annual terms) over the period for

Figure 8 Optimal Inflation, Trading Off Inflation Target and Price-Adjustment Costs

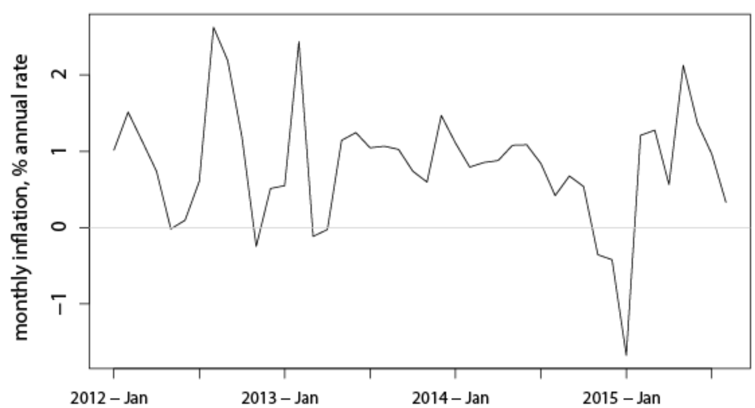

which the Federal Reserve has had a formal inflation target. The lines plotted in these figures are simply the average of 2 percent and the lines plotted in Figures 5 and 6. Thus, for the most recent period, instead of optimal deflation at around 1.25 percent, there is optimal inflation at around 0.4 percent. Actual 12-month inflation has been between 0.15 percent and 0.35 percent throughout 2015.

Inflation-Targeting Era: Trade-off When Adjustment Costs are Relative to 2 Percent Price Increase

The fact that there is little long-run trend in relative energy price means that the measure of optimal inflation in Figures 8 and 9 is biased away from the Fed's official 2 percent inflation target. One might find this property problematic: Given that the Federal Open Market Committee has stated its intention to keep inflation around 2 percent, there are limits to the immediate practical usefulness of a prescription that the Fed keep average inflation below 2 percent. To address that concern, here we modify the optimal inflation measure from Figures 8 and 9 by positing that the costs of price adjustment are incurred relative to a 2 percent change in prices rather than relative to no change in prices.

Figure 9 Optimal Monthly Inflation, Measured Over 12 Months, Trading Off Inflation Target and Price-Adjustment Costs

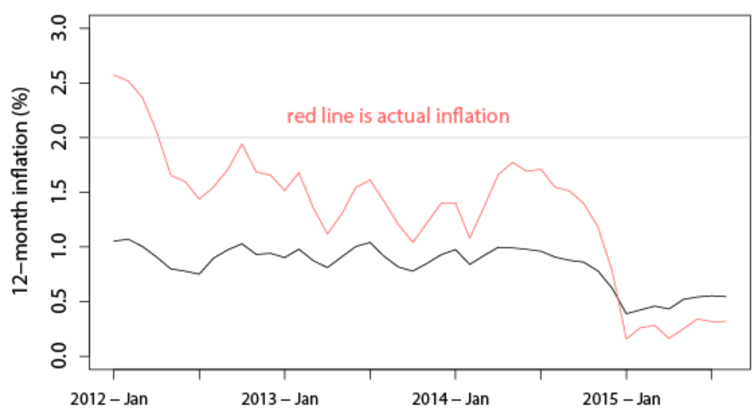

That is, instead of (15) we use

$$\xi_{k,t}(z) = \frac{\theta_k}{2}\left(\frac{y_{k,t}(z)}{a_{k,t}}\right)\left(\frac{P_{k,t}(z)}{P_{k,t-1}(z)} - (1.02)^{1/12}\right)^2. \tag{34}$$

With this modification, the inflation rate that minimizes adjustment costs will be 2 percent on average unless there is a trend in relative prices. The assumption behind (34) is not entirely ad hoc. In a world with expectations anchored at 2 percent inflation, it seems plausible that households and firms would adapt so that there would eventually be negligible costs associated with prices changing at a 2 percent rate.

Figures 10 and 11 plot the modified-optimal inflation rate that is a simple average of 2 percent and the inflation rate that minimizes sector 2 price-adjustment costs when those costs are incurred relative to a 2 percent trend. Focusing on Figure 11, the 12-month inflation rate, two features stand out. First, as expected, the optimal inflation rate in 2015 rises relative to what the previous figures showed; it is currently around 1.5 percent, which is well above the actual inflation rate, indicated by the red line in the figure. Second, although the fluctuations have been small, since 2012 the optimal inflation rate has been persistently below 2 percent, even though both components of the calculation tie that rate

Figure 10 Optimal Inflation, Trading Off 2 Percent Target and Price-Adjustment Costs Relative to 2 Percent

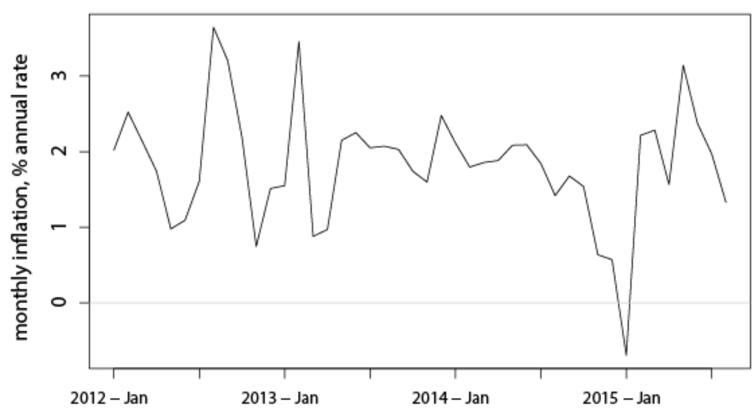

to 2 percent over the long run. The obvious explanation is that since 2012 there have been repeated large declines in relative energy prices.

For Figure 12, we plot a second version of optimal inflation that is centered around 2 percent. In this case the optimality criterion puts full weight on minimizing adjustment costs relative to a 2 percent rate of price change. Because no weight is placed on achieving the 2 percent inflation target in a given period, optimal inflation declines more in early 2015 than when the 2 percent inflation target receives equal weight, in Figure 11. It is still the case, however, that the optimal inflation rate in Figure 11 has been above actual inflation in recent years.

5. CONCLUSION AND CAVEATS

While monetary policy's goals are typically framed in terms of aggregate inflation, the aggregate inflation rate is an outcome of many individual price changes. Relative prices across goods and sectors undergo long-run trend changes and dramatic monthly fluctuations. How those factors affect the optimal inflation rate has been our focus. With respect to long-run trends, we used a particular model with costly price adjustment to show how the optimal inflation rate can vary with

Figure 11 Optimal Monthly Inflation, Measured Over 12 Months, Trading Off 2 Percent Target and Adjustment Costs Relative to 2 Percent

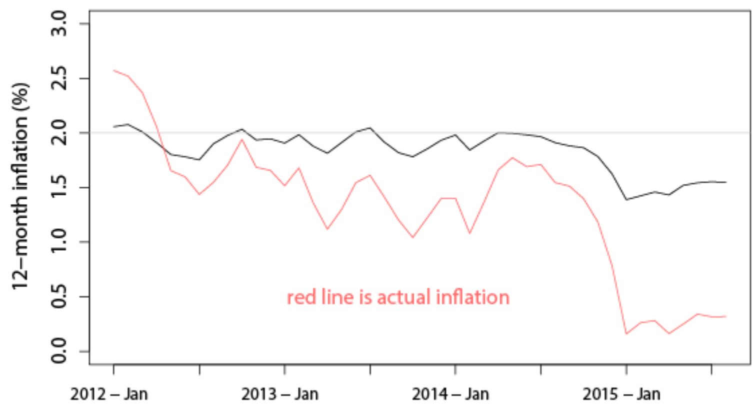

relative productivity growth across sectors and relative price-adjustment costs. With respect to short-run fluctuations, we narrowed the analysis to the behavior of relative energy prices in the United States, tracing out measures of the inflation rate over time that would have minimized price-adjustment costs according to the model. The stark nature of the model means that the results on steady-state inflation should be viewed as suggestive rather than definitive. However, the qualitative result that deflation (inflation) is optimal when price stickiness is relatively high in sectors that have relatively low (high) productivity growth is somewhat general, having also been shown to hold in models of staggered price setting and fixed costs of price adjustment (Wolman 2011). The results on period-by-period inflation may be of more practical use because they do not rely on the entire model and they emphasize the relative price behavior of the most volatile component of the consumption price index, energy goods and services. In line with Aoki (2001), we found that the dramatic decline in energy prices in 2015 was associated with a marked decline in the optimal inflation rate, even when the optimality criterion put a weight of one-half on meeting the 2 percent inflation target.

Figure 12 Optimal Monthly Inflation, Measured Over 12 Months, Adjustment Costs Relative to 2 Percent

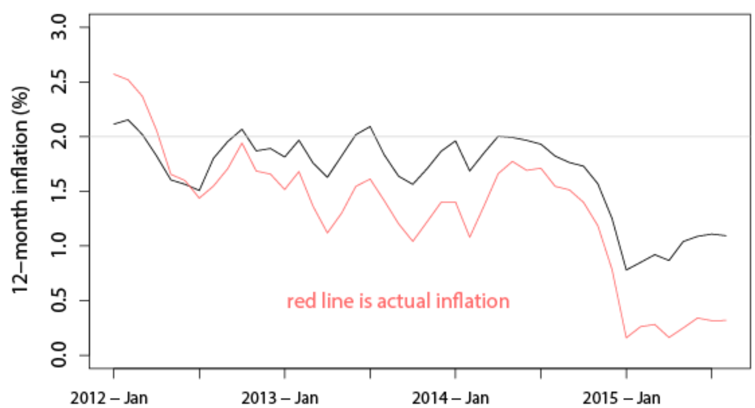

Throughout the paper, relative price changes were viewed as exogenous; in the steady-state analysis relative price changes were generated by exogenous differences in productivity growth across sectors, whereas in the short-run analysis relative price changes themselves were taken directly as inputs into the determination of optimal policy. Of course, any mechanism for monetary nonneutrality will generally lead to relative price changes being in part endogenous with respect to monetary policy. Given our ignorance about the details of that endogeneity however, proceeding under the assumption that relative price changes are exogenous seems like a reasonable way to proceed, although certainly not the only reasonable way. For simplicity, and in order to focus attention on the behavior of energy prices, the model in this paper had only two sectors, whereas the number of consumption categories measured by the Bureau of Economic Analysis is in the hundreds. In principle, one could perform the same kind of analysis in a model with many more sectors.

Throughout the paper, we assumed that monetary policy could directly control the inflation rate. For analyzing long-run trends this assumption seems appropriate, but in the shorter run it is clearly problematic. Any discussion of monetary policy and the behavior of inflation in 2014 and 2015 needs to confront the question of how accurately and over what horizon a central bank can control inflation. Here it was simply assumed that the central bank has perfect control, and optimal inflation was derived under that assumption. Perfect control is an important benchmark. With imperfect control, large relative price shocks will likely lead to greater fluctuations in the optimal inflation rate. That perspective suggests that optimal inflation in 2015 may have been below the levels in Figures 11 and 12. This is clearly a topic on which further work is needed.

APPENDIX: STEADY-STATE WELFARE EXPRESSION

In this Appendix we derive the closed-form solution for steady-state welfare, used to produce Figures 1–4. From (25)–(29), we have

$$p_{1,0} = \frac{\varepsilon}{(\varepsilon-1)(1+\tau)} w_0 \Gamma_1(\pi) \tag{35}$$

$$p_{2,0} = \frac{\varepsilon}{(\varepsilon-1)(1+\tau)} w_0 \Gamma_2(\pi) \tag{36}$$

$$\frac{p_{2,0}}{1-\nu} = \left(\frac{1-\nu}{\nu}\frac{p_{1,0}}{p_{2,0}}\right)^{-\nu}, \tag{37}$$

where

$$\Gamma_1(\pi) = 1-\theta_1\left(\left(\frac{\mu_1}{\mu_2}\right)^{\nu-1}\pi-1\right)\times \left[\left(\frac{1-\beta}{\varepsilon}\right)\left(\frac{\mu_1}{\mu_2}\right)^{\nu-1}\pi-\frac{1}{2}\left(\left(\frac{\mu_1}{\mu_2}\right)^{\nu-1}\pi-1\right)\right]$$

and

$$\Gamma_2(\pi) = 1-\theta_2\left(\left(\frac{\mu_1}{\mu_2}\right)^{-\nu}\pi-1\right)\times \left[\left(\frac{1-\beta}{\varepsilon}\right)\left(\frac{\mu_1}{\mu_2}\right)^{-\nu}\pi-\frac{1}{2}\left(\left(\frac{\mu_1}{\mu_2}\right)^{-\nu}\pi-1\right)\right].$$

It follows that

$$\frac{p_{1,0}}{p_{2,0}} = \Gamma_1(\pi)/\Gamma_2(\pi). \tag{38}$$

We can then solve for $p_{1,0}, p_{2,0}$, and w_0 :

$$p_{2,0}(\pi) = (1-\nu)\left(\frac{1-\nu}{\nu}\Gamma_1(\pi)/\Gamma_2(\pi)\right)^{-\nu}, \tag{39}$$

$$w_0(\pi) = \frac{(\varepsilon - 1)(1 + \tau)}{\varepsilon} \left(\frac{\Gamma_1(\pi)}{\nu} \right)^{-\nu} \left(\frac{\Gamma_2(\pi)}{1 - \nu} \right)^{\nu - 1}, \tag{40}$$

and

$$p_{1,0}(\pi) = (1 - \nu) \left(\frac{1 - \nu}{\nu} \right)^{-\nu} \left(\Gamma_1(\pi) / \Gamma_2(\pi) \right)^{1 - \nu}. \tag{41}$$

Consumption follows from (4) and (5):

$$c_0(\pi) = \frac{(\varepsilon - 1)(1 + \tau)}{\varepsilon \chi} \left(\frac{\Gamma_1(\pi)}{\nu} \right)^{-\nu} \left(\frac{\Gamma_2(\pi)}{1 - \nu} \right)^{\nu - 1}.$$

The last variable we need in order to compute welfare is labor input. From (18), we have

$$h_0 = \frac{y_{1,0}}{a_{1,0}} \cdot \left(1 + \frac{\theta_1}{2} (\pi_1 - 1)^2 \right) + \frac{y_{2,0}}{a_{2,0}} \cdot \left(1 + \frac{\theta_2}{2} (\pi_{2,0} - 1)^2 \right). \tag{42}$$

This expression requires that we know $y_{1,0}$ and $y_{2,0}$. From (21) and (22), we can easily compute $y_{1,0}$ and $y_{2,0}$ using our solutions for the wage and relative prices:

$$y_{1,0}(\pi) = \frac{\nu}{\chi} \frac{w_0(\pi)}{p_{1,0}(\pi)}, \tag{43}$$

$$y_{2,0}(\pi) = \frac{1 - \nu}{\chi} \frac{w_0(\pi)}{p_{2,0}(\pi)}. \tag{44}$$

Then using (42) to compute labor input, we can evaluate welfare as a function of inflation:

$$\begin{aligned} W(\pi) &= \sum_{t=0}^{\infty} \beta^t (\ln c_t - \chi h_0) \\ &= \sum_{t=0}^{\infty} \beta^t \left(\ln c_0(\pi) + t \ln \left(\mu_1^{\nu} \mu_2^{1-\nu} \right) - \chi h_0(\pi) \right). \end{aligned}$$

REFERENCES

Aoki, Kosuke. 2001. "Optimal Monetary Policy Responses to Relative-Price Changes." *Journal of Monetary Economics* 48 (August): 55–80.

Bils, Mark, and Peter J. Klenow. 2004. "On the Importance of Sticky Prices." *Journal of Political Economy* 112 (October): 947–85.

Bodenstein, Martin, Christopher J. Erceg, and Luca Guerrieri. 2008. "Optimal Monetary Policy with Distinct Core and Headline Inflation Rates." *Journal of Monetary Economics* 55 (October): S18–S33.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (September): 383–98.

Friedman, Milton. 1969. "The Optimum Quantity of Money." In *The Optimum Quantity of Money and Other Essays,* edited by Milton Friedman. Chicago: Aldine Publishing Company, 2–50.

Levy, Daniel, Mark Bergen, Shantanu Dutta, and Robert Venable. 1997. "The Magnitude of Menu Costs: Direct Evidence From Large U.S. Supermarket Chains." *Quarterly Journal of Economics* 112 (August): 791–825.

Rotemberg, Julio J. 1982. "Sticky Prices in the United States." *Journal of Political Economy* 90 (December): 1187–211.

Wolman, Alexander L. 2008. "Nominal Frictions, Relative Price Adjustment, and the Limits to Monetary Policy." Federal Reserve Bank of Richmond *Economic Quarterly* 94 (Summer): 219–33.

Wolman, Alexander L. 2011. "The Optimal Rate of Inflation with Trending Relative Prices." *Journal of Money, Credit and Banking* 43 (March-April): 355–84.