Monetary Policy Behind the Veil of Ignorance

Rommeswinkel, Hendrik; Wu, Hao*

February 12, 2020

Abstract

We analyze the problem of the choice of a central bank constitution. We model the decision problem as a choice behind a veil of ignorance in which the policy maker only receives information about predicted behavior under different policies. The policy maker is informed about (probability distributions of) consumers' behavior and the distribution of productivity shocks but does not know consumers' interpersonally comparable utility functions. Starting from a representation theorem for the policy maker's preferences over policies, we compare price stabilization, output stabilization, and inflation targeting in a standard new Keynesian model with Calvo price staggering. Surprisingly, under our policy criterion, the policy maker perceives a tradeoff between output stabilization and price stabilization. The reason is that in the absence of knowledge about cardinal utility functions, stabilizing the natural level of output is not normatively desirable. We find that the policy maker puts a higher emphasis on price stability than output stability if price staggering is low, intertemporal discounting is high, intertemporal substitutability is low, or substitutability between goods is high.

Keywords: Monetary Policy, Monetary Constitution, Veil of Ignorance, Robustness, Freedom of Choice, Policy Criteria

JEL Classification: E12, E52, E61

*Department of Economics, National Taiwan University, No. 1, Sec. 4, Roosevelt Rd., Taipei 106.

1 Introduction

Utilitarian economic policy evaluation requires not only an economic model of how policies affect individual behavior but also cardinal information about individual's utilities. However, cardinal, interpersonally comparable utility functions are notoriously difficult to obtain.[1] Undergraduate students are thus often cautioned not to interpret utility functions in consumer theory as cardinally meaningful. The difficulty of obtaining cardinal utility information has led economists to instead employ less stringent criteria such as Pareto efficiency. This however often leads to a large set of optimal policies.

The problem of using utility information of either ordinal or cardinal scale is amplified by the insights gained from behavioral economics – if agents do not maximize a well-defined utility function, the policy maker runs the risk of maximizing the wrong objective. In this paper we address these problems in the context of monetary policy by examining what monetary policy a rational policy maker would choose if supplied with a model of the behavior of agents but without information about the utility scales of the agents. The absence of useful utility information is especially plausible in the context of the choice of a monetary constitution. Basing a monetary constitution not only on the stability of aggregate behavior over time but also on the stability of individuals' imputed utility functions is highly error-prone. A monetary constitution that is robust to misspecification of the utility functions is therefore highly desirable.

It is a recurring theme in economics to create institutions such that the policy maker does not need to know the exact utility function of individuals. For example, the first welfare theorem gives conditions under which a market may reach an efficient allocation without intervention of the policy maker. Another example is the mechanism design literature that commonly designs institutions in such a way that for some set of utility types agents voluntarily reveal their type and reach the desired outcomes. In such settings, instead of implementing a particular outcome, the policy maker prefers to leave decision room for the agent. Our policy maker's decision criterion reflects this idea by maximizing a measure of control the agents have over their outcomes. Thus, in the absence of utility information the policy maker maximizes the control the consumers have over their own consumption as measured by the informativeness of the demand functions about consumption outcomes.

We justify the criterion of the policy maker axiomatically using a decision

[1] For an overview of the literature, see Elster and Roemer (1991).

problem similar to that of Harsanyi (1977). Behind a veil of ignorance, the policy maker faces uncertainty about the future behavior of consumers and external shocks. We impose the von Neumann-Morgenstern axioms on the policy maker for uncertainty generated purely by external shocks. That is, if two policies are effectively lotteries over outcomes in which the consumers have no influence, the policy maker compares the policies by their expected desirability. In addition, we impose that policies that are observationally indistinguishable must be indifferent to each other. In other words, if a monetary policy yields exactly the same predicted behavior as another monetary policy, then the policy maker is indifferent between the two policies. This captures the idea that economic models are only useful as descriptions of (observable) behavior and any model-imposed unobservable quantities (such as utility) are normatively irrelevant beyond what can be deduced from behavior.[2] Two further conditions guarantee additive separability of the policy criterion. First, we impose additive separability of policies into subpolicies that only affect a subset of the outcomes. Second, we impose additive separability across choices made by the agents. If a change in policy makes the choice between two actions irrelevant, then the entailed loss in value is independent from the other actions taken. From these axioms follows that the decision maker ranks policies by the mutual information between actions and outcomes and an expected valuation of the outcomes. Thus, the policy maker maximizes the expected control of consumers over consumption, which we interpret as a measure of freedom of choice.[3]

We model the information of the policy maker about the economy using a standard new Keynesian model with sticky prices. We employ a discrete time model of an economy in which a representative agent has constant elasticity of substitution (CES) preferences over a continuum of goods. The marginal rate of substitution of each of these goods is stochastic behind the veil of ignorance; preferences over goods are unknown at the time of policy choice and may vary over time. The agent provides labor to a continuum of firms that each produce a single good. Firms' productivity is also stochastic over time. There are both firm-specific shocks and shocks to the productivity of all firms. We employ Calvo (1983) price staggering; every period, only a fraction $1 - \alpha$ of firms get to reset their prices, and all other firms must sell their goods at the same price as

[2]This is indeed our only axiom that is not fulfilled by an expected aggregate utility maximizer. A utilitarian policy maker would make the evaluation of policies not only dependent on observable behavior but also on the utility scales of the agents.

[3]For an overview over the freedom of choice literature, see Dowding and van Hees (2009).

before. Firms are monopolists and maximize profits. In this setting, monetary policy affects real variables.

In the standard new Keynesian monetary model there is usually no tradeoff between fully stabilizing prices or the output gap. This “Divine Coincidence” suggests that adopting a monetary policy constitution that prescribes price stability or output stability yields equivalent results. Blanchard and Galí (2007) showed that under real imperfections this result breaks down. We extend their result by showing that if the policy maker is uninformed about utility scales, then even in a model *without* real imperfections the policy maker perceives a tradeoff between output and price stabilization. The reason is that if the standard new Keynesian model is interpreted as only describing behavior and the utility functions are not taken literally as the exact utility functions of the agents, then the natural level of output inferred from the model is not necessarily the most desirable (based on the unknown true utility of the agents). Since our policy maker cannot know the natural level of output, this creates a tradeoff between *absolute* output level stabilization and price stabilization.

Under Calvo price staggering, a utilitarian policy maker faces every period a tradeoff between inefficient consumption of goods with rigid prices (relative to labor) and inefficient consumption of goods with flexible prices (relative to labor). To optimally solve this tradeoff, knowing the correct utility function is necessary. Rotemberg and Woodford (1999) showed that under second order approximations of utility, a policy maker solves this tradeoff by minimizing a loss function that expresses a tradeoff between stabilizing the output gap, stabilizing inflation, and minimizing inflation. Our policy maker instead tries to minimize the degree to which external shocks and price fluctuations influence consumption. A policy enforcing full price stability guarantees that all goods are affected equally by aggregate shocks. However, aggregate shocks fully affect consumption. The downside of dampening aggregate shocks is that consumption of a good now depends on whether its price is flexible in this time period or not. The policy maker therefore faces the tradeoff between reducing external shocks and the disturbances in the consumption associated with this reduction.

We analyze the choice between a policy in which the price level is kept constant, a policy in which aggregate output is kept constant, and inflation targeting. We obtain a criterion under which price stabilization is superior to output stabilization. Low intertemporal elasticity of substitution of leisure and consumption makes price stabilization more attractive. Similarly, a high elasticity of substitution between different products makes price stabilization more attractive. A higher price rigidity and lower intertemporal discounting (i.e., a higher discount factor) makes output stabilization more attractive. These results are confirmed in our analysis of inflation targeting; the same comparative statics hold with respect to the intensity to which the central bank should react to the deviation from the inflation target.

Our optimal policy results are *robust* in the following sense. If two economists disagree on whether a utility-maximizing model of consumers or a behavioral model is correct but both models yield the same predicted behavior, then the policy maker will have the exact same preferences over policies. Thus, the policy maker only requires a descriptively accurate model of behavior, not a literally accurate model. Whether consumers *actually* maximize a well-defined utility function is irrelevant to the policy maker as long as the model accurately describes behavior. We consider this robustness a desirable feature of our criterion.

The paper continues as follows. In Section 3, we axiomatically derive mutual information between demand and consumption outcomes as the policy maker's objective. In Section 4, we introduce the model which represents the information received by the policy maker about the effects of the policy. In Section 5, we derive the freedom of choice obtained from various policies and compare under which circumstances one or another policy is more desirable. Section 6 presents avenues for further research and concludes.

2 Literature

Monetary policy at different time periods and with different emphasis of topics is surveyed by Blanchard, Dell'Ariccia, and Mauro (2010), Blinder, Ehrmann, Fratzscher, De Haan, and Jansen (2008), Clarida, Gali, and Gertler (1999), Flood and Isard (1989), Friedman (1988), Goodfriend (2007), Taylor (1999). Following Rotemberg and Woodford (1997, 1999), utility as a welfare criterion in monetary policy has been explored in many directions. Most contributions centered around refining the economic model of behavior. Debortoli, Kim, Lindé, and Nunes (2019) discusses how to design simple loss functions for monetary policies which is closest in spirit to our paper. Our analysis suggests a loss function that is motivated by (in principle) empirically accessible information only.

There is a large literature on macroeconomics with boundedly rational agents. For surveys, see for example Akerlof (2002), Rötheli (2015), Shiller (2003). Usually, behavioral insights are incorporated in order to improve the predictions of models. Our approach is instead to ensure that policies remain normatively convincing even if the policy maker is uncertain about the rationality of agents. Given that our policy maker maximizes an information-theoretic measure, there is also an interesting connection to the literature on rational inattention (Sims, 2003), in which these information-theoretic measures arise within the utility functions of agents. At this point speculative –but interesting– is the idea that combining such models with an information-theoretic objective of a policy maker makes it easier to obtain analytical results for welfare as the consumers' and policy makers' objectives are more aligned.

Parameter uncertainty (Edge, Laubach, & Williams, 2010; Sala, Söderström, & Trigari, 2008) and model uncertainty (Coenen, 2007; Giannoni, 2002; Levin & Williams, 2003) are closely related to the problem we address. In such models, the policy maker faces uncertainty about the behavioral model but is certain about the utility scales for each possible model of behavior. We make the assumption that the policy maker cannot even form beliefs over a parametric family of utility functions. In contrast, the policy maker has an exact model of the behavior of consumers.

Uncertainty about the natural rate of interest (and thus the other "natural" variables) has been historically much discussed in the literature. Orphanides (2003a, 2003b), Orphanides and Williams (2002) give an overview of the empirical difficulties of dealing with noisy estimates of these "natural" variables in monetary policies. The Austrian school already very early addressed the problem of optimal monetary constitutions (Boettke & Smith, 2016; D'Amico, 2007, provide surveys).

Our analysis assumes that the monetary constitution can constrain the central banker effectively to avoid time-consistency problems of the form discussed by (Barro & Gordon, 1983; Kydland & Prescott, 1977). A monetary policy constitution of course sets not only a policy goal but also creates institutions that implement the policy. The analyses of Lohmann (1992) and Walsh (1995) complement ours in this respect.

Within welfare economics, our analysis is closely related to impartial observer theorems Harsanyi (1953, 1955, 1977) and the Harsanyi-Sen debate on utilitarianism as a policy criterion (Grant, Kajii, Polak, & Safra, 2010; Sen, 1977). Weymark (2011) provides a survey on this debate. Our suggested criterion for the policy maker can be interpreted as a measure of freedom of choice. For the axiomatic derivation of the measure in a game theoretic context, see Rommeswinkel (2019). Dowding and van Hees (2009) surveys the freedom of choice literature.

3 Normative Framework

The main normative framework in economics is welfarism, which aggregates individual preferences into society's preferences via the concept of Pareto efficiency; a state of the world is better than another, if all individuals agree that the former is at least as good as the latter and at least one individual strictly prefers the former to the latter. Many remarkable results can be derived from this rule but its simplicity comes at a cost; it usually does not yield a complete ordering of policies but only a partial order. In a series of papers, Harsanyi (1953, 1955, 1977) provided a foundation for a criterion that yields a complete order, utilitarianism. Harsanyi (1977) assumed that the policy maker has to make a decision behind the veil of ignorance about which position in society she will occupy. The policy maker is an expected utility maximizer and the individuals in each position in society are expected utility maximizers. If the policy maker subscribes to reduction of compound lotteries, then the policy maker should maximize the weighted sum of the expected utility functions of the positions in society. However, this requires that the policy maker knows the expected utility functions associated with each of the possible positions in society. Most importantly, the scale of the utility function must be known for every individual and be interpersonally comparable to other individuals. Scepticism whether this strong requirement can be fulfilled is warranted. Even if we were to obtain lottery choice data for all individuals, this lottery data were consistent with expected utility maximization, and we were able to estimate the shape of the utility function from the data, still these functions would only be unique up to separate affine transformations for each individual. Moreover, as argued by Sen (1977), Grant et al. (2010), Weymark (2011), it may still be normatively more compelling to aggregate individual utilities using a nonlinear aggregation function.

We therefore return to Harsanyi's initial setting but do not allow the policy maker to form preferences dependent on unobservable data such as utility scales – instead, the policy maker must rank policies solely on behavior data. We assume that the behavior data is given by the demand functions of the individuals. Behind the veil of ignorance, the policy maker is uncertain about the individual demand functions and the production possibilities of the economy. The policy maker has to decide behind the veil of ignorance what kind of monetary constitution to adopt. Different monetary policies may lead to different behavior and thus different outcomes for consumers.

The problem of choosing a central bank constitution closely resembles the assumptions made by the veil of ignorance framework we employ. Central banks are usually highly independent institutions bound by certain constitutional constraints. When choosing the constitutional constraints, policy makers face high uncertainty about future shocks to demand and supply. It is therefore unrealistic to assume that a policy maker will be able to perform a meaningful utilitarian analysis in which utility scales are known. Instead, we only impose a set of axioms on the policy maker's preference over the set of policies $\mathcal{P}$. The set of policies $\mathcal{P}$ is isomorphic to a set of probability measures – each policy $P \in \mathcal{P}$ can be represented by a probability measure that describes the expectations of the policy maker over the behavior of the individuals under the policy P. For simplicity, we denote both a policy and its corresponding probability measure by the letter P. In the context of our policy maker's problem, the probability measure P represents the policy maker's information about how agents (including consumer, firms, and nature) interact in a market given a fixed monetary policy.

We assume that there exists a set of outcomes $\mathcal{O}$ that consist of the quantities $c \in \mathbb{R}_{+}^{I}$ of the goods indexed by I consumed by the agent and the hours worked by the agent, $y \in \mathbb{R}$. A generic outcome is denoted by the letter $o \in \mathcal{O}$. The agent reports a demand function $x : (p, w) \mapsto (c, y)$ that maps prices and wages into outcomes subject to the constraint $pc \equiv \int_{i \in I} p_i c_i = yw$.[4] We denote the set of demand functions of the consumer by $\mathcal{X} = \{(p, w) \mapsto (c, y) : pc = yw\}$, i.e., the set of all functions that map prices into quantities subject to the budget constraint. After the consumer reports the demand function, the Walrasian auctioneer allocates consumption. Since there may exist uncertainty about

[4]We derive the policy maker's preference in a model with only a decision in a single time period. This is without much loss of generality; in the intertemporal decision problem, we could replace the demand functions by demand functions over time and consumption outcomes by sequences of consumption quantities.

the production conditions, etc., from the position of the policy maker every policy yields only a probabilistic relationship between the demand and the outcome. Under the assumption that policies differ meaningfully if and only if the resulting behavior or outcomes differ, the set of policies can be assumed to be equal to the set $\mathcal{P}$ of finite support probability measures on $(\mathcal{X} \times \mathcal{O})$ endowed with the product topology.[5]

The policy maker forms preferences over the set of policies $\mathcal{P}$. A standard normative assumption on the policy maker is rationality:

Axiom 1 (Rationality)**.** The policy maker forms complete and transitive preferences $\succsim$ over the set of policies $\mathcal{P}$.

Continuity is also a standard assumption that guarantees that the policy maker ranks similar policies similarly.

Axiom 2 (Continuity)**.** The policy maker's preference is continuous.

Continuity guarantees that there are no "jumps" in the evaluation of policies in case behavior or outcome probabilities change by a little.

We define policies policies in which the choices of the consumers do not matter as lotteries. That is, a policy is a lottery if every demand function in the support of P yields the same conditional probability distribution over consumption outcomes. In other words, in such policies the Walrasian auctioneer completely ignores the stated demand function and instead randomly assigns consumption to the consumers according to the marginal distribution $P[o]$.

For lotteries, it is natural to assume the von Neumann-Morgenstern independence axiom. Define $P''' = \alpha P \oplus (1-\alpha)P''$ as the probability mixture such that $P'''[x,c] = \alpha P[x,c] + (1-\alpha)P''[x,c]$.

Axiom 3 (Lottery Independence)**.** For all policies that are lotteries, the policy maker obeys the von Neumann-Morgenstern Independence axiom, i.e., if P, P', and P'' are lotteries, then $P \succsim P'$ if and only if $\alpha P \oplus (1-\alpha)P'' \succsim \alpha P \oplus (1-\alpha)P''$ for all $\alpha \in (0,1)$.

The von Neumann-Morgenstern independence axiom is a standard axiom of rationality for decisions under risk. Note that in contrast to Harsanyi (1977), we only impose expected utility rationality on the policy maker (which seems

[5]For simplicity, we introduce the following notational conventions. $P[x,o] \equiv P[\{(x,o)\}]$, i.e., for singletons we omit the set notation. Marginal probabilities are defined as $P[o] \equiv P[\{(x',o') : o' = o\}] = \int_{x' \in X} P[x,o]$. The conditional measure uses the $|$ notation, $P[o|x] = \frac{P[o,x]}{P[x]}$ if $P[x] > 0$.

obviously desirable) but not on those affected by the policies. The consumers may in principle follow any behavioral model according to the information of the policy maker. Moreover, we only impose the independence axiom on the policy maker in case the consumers have no meaningful choices. If a policy gives consumers some way of influencing the outcome with their reported demand, it is not a lottery and lottery independence does not impose restrictions on the policy maker's preferences. This allows the policy maker to treat uncertainty derived from consumer behavior different from uncertainty from external shocks. For example, suppose the policy maker is indifferent between policies that dictatorially assign either outcome o or o'. Both policies are (trivial) lotteries with $P[o] = 1$ and $P'[o'] = 1$. Lottery Independence imposes that $P \sim P'$ implies that the policy maker is also indifferent to any policy that randomly assigns either o or o' with probabilities α and $1-\alpha$, $P \sim \alpha P \oplus (1-\alpha)P'$. However, we permit that the policy maker may strictly prefer a policy in which the consumer *chooses* either o or o' with probability α and $1-\alpha$, respectively. Formally, if $P''[o|x] = 1$ and $P''[o'|x'] = 1$, and $P''[x] = \alpha = 1 - P[x']$, then $P'' \succ P' \sim P$ is an admissible preference. This is consistent with the idea that the policy maker may value control of the consumer over outcomes.

For a given policy, two demand functions are indistinguishable in case they yield the same conditional distribution of consumption. When the policy maker obtains data about behavior, we assume that all the policy maker can observe is how likely it is that somebody will obtain a certain consumption level given a reported demand function. For example, two demand functions derived from utility functions that are monotone transformations of another are observationally equivalent. Similarly, two demand functions that only differ on unavailable goods (for example, goods that are not produced under policy P) are also observationally equivalent. Formally, in a policy P, two demand functions x and x' are equivalent if $P[o|x] = P[o|x']$ for all o. Two policies P, P' are observationally equivalent, denoted by $P \approx P'$, if all the matching equivalence classes of demand functions are equally likely.

Axiom 4 (Observational Equivalence)**.** If two policies are observationally equivalent, then the policy maker is indifferent between the policies.

We illustrate the axiom with an example. Suppose a policy maker prohibits the consumption of two goods, drugs and rock'n'roll. Consider the demand function x of a drug addict and the demand function x' of a rock'n'roll addict. Suppose that the two demand functions only differ with respect to the consumption of the two goods (whenever these goods are available). Then, under the policy in which the policy maker prohibits the consumption of both goods, the two demand functions are observationally equivalent. This is because we have empirically no means to distinguish drug addicts from rock'n'roll addicts – both behave the same way given the prohibition of drugs and rock'n'roll. According to the Observational Equivalence Axiom, the relative likelihood of a consumer being a drug addict instead of a rock'n'roll addict $P[x]/P[x']$ does not matter for the preference of the policy maker for this policy. Indifference of two policies directly follows if they both prohibit drugs and rock'n'roll and only differ on the relative likelihood $P[x]/P[x']$, but agree on $P[x] + P[x']$ and $P[x'']$ for all other demand functions x''. It is straightforward to construct similar examples in which behavior is observationally equivalent because of different utility scales, randomized allocation of goods, or changes of the underlying behavioral model.

We next give a formal structure to the idea that in some policies, consumers have more influence on their outcomes than in other policies. From this, we then derive an axiom that captures the intuition that the policy maker's preference is independent across changes in influence over different demands.

Suppose that the policy maker expects that under a certain policy P, the consumer reports demand function x or x' to the Walrasian auctioneer. The two demand functions are initially observationally distinct. For example, if a fixed share of the budget is reserved for goods i and i', then x might allocate the entire share to i while x' allocates the entire share to i'. A choice deprivation is a change in the policy such that the Walrasian auctioneer ignores the distinction between x and x' and randomly allocates goods either according to x or x', ignoring the actually reported demand.

We assume now that the policy maker's preference is (ceteris paribus) independent in identical choice deprivations. Consider two policies that differ only on the joint probability of outcomes and the demand functions in $\mathcal{X}\backslash\mathcal{X}'$. Assume that in all other respects, the policies are identical. Due to their differences, one policy may of course be preferable to the other. We assume that this preference does not change if the policy maker deprives the consumer of demand choices in $\mathcal{X}'$. That is, if the policy maker deprives the consumer in two policies of identical choices, then the preference between the two policies does not change. Formally, we assume the following axiom.

Axiom 5 (Deprivation Independence)**.** The preferences of the policy maker are independent of identical choice deprivations; applying an identical choice deprivation to two policies does not change the preference between these policies.

We define a sub-policy as the policy obtained from conditioning the probability distribution P to a subset of the consumption outcomes. Formally, if P is a policy and $\mathcal{O}'$ is a subset of outcomes $\mathcal{O}$, then the probability measure obtained by conditioning P on $\mathcal{O}'$ is a subpolicy of P with respect to $\mathcal{O}'$

Axiom 6 (Weak Decomposability)**.** If two policies differ only on their subpolicies with respect to the outcomes $\mathcal{O}'$, then the preference between the two policies is determined by the preference on the subpolicies with respect to $\mathcal{O}'$.

Weak Decomposability states that policies can be improved by focusing on improvements on a subset of the outcomes. This is an assumption that is implicit in much of economic analysis that finds improvements in localized contexts to gain overall improvements. For example, in general equilibrium analysis an intervention in some market can be separately analyzed from other markets if it does not influence behavior in the other markets. We call this condition weak because it is only required to hold if outside of the subpolicy all behavior remains observationally identical; if a change in a subpolicy were to affect behavior outside of $\mathcal{O}'$, improvements of that subpolicy do not necessarily yield improvements of the overall policy.

Theorem 1. *If the policy maker's preference fulfills Axioms 1-6, then the policy maker's preferences $\succsim$ can be represented by a function of the form*

$$U[P] = \sum_{o\in\mathcal{O}} P[o]v[o] + r \cdot \sum_{x\in\mathcal{X}} \sum_{o\in\mathcal{O}} P[x,o] \ln \frac{P[x,o]}{P[x]P[o]} \tag{1}$$

where all $v[o]$ and r are real valued parameters to the policy maker's utility representation.

The first component of the criterion is an expectation over the policy maker's valuation of the outcomes. This can be seen as the policy maker's instrumental valuation of choosing some policy over another. The second component is the mutual information between demand and outcomes. This component will in the following be called a measure of freedom of choice. The interpretation is that in the absence of utility information, the policy maker maximizes the degree to which individuals control their outcomes.

In our analysis of monetary policy, we assume on the domain of policies we consider, $v[o]$ is constant and without loss of generality equal to zero. We assume this for the following reason: $v[o]$ can be interpreted as the utility function of the policy maker over the outcomes. Normatively, maximizing the expectation of $v[o]$ is only interesting if $v[o]$ is some heuristic to measure consumer welfare, for example income. Sensible policies that combine the maximization of the expected $v[o]$ with the maximization of mutual information would therefore simply represent an intermediate case between maximizing a simple welfare measure and the analysis we perform below.

The policy maker's criterion reduces to:

$$U[P] = I[x;o] = \sum_{x\in\mathcal{X}} \sum_{o\in\mathcal{O}} P[x,o] \ln \frac{P[x,o]}{P[x]P[o]} \tag{2}$$

This is the mutual information between demand and outcomes. Mutual information is a measure of correlation that assumes very little about the functional relationship between the variables. It therefore measures statistical dependence not only in case of linear relationships (as the correlation coefficient does). However, in case of linear relationships of jointly normal distributed variables, mutual information is ordinally equivalent to the correlation coefficient. Moreover, we are primarily interested in the degree to which consumers control their consumption outcomes in our model. We therefore assume that two outcomes are distinct if and only if any of the consumed quantities are distinct. Since we deal with real valued quantities we replace the measure P by a joint density and the summation by integration. While the axiomatization holds for finite support probability measures, the generalization to an infinite support is technically nontrivial but conceptually straightforward. The mutual information for a joint density P over consumption c and an arbitrary[6] parametrization of the demand functions is given by:

$$I(x;c) = \int_x P(x) \int_c P(c|x) \ln \frac{P(x,c)}{P(x)P(c)} \tag{3}$$

4 The Model

The policy evaluation naturally depends on the information of the policy maker about the impact of the policy on the economy. We assume the policy maker's

[6] Mutual information is invariant under homeomorphisms, see Kraskov, Stögbauer, and Grassberger (2004).

information is given by a DSGE model with uncertainty about the agent's preferences and about the productivity of firms. The policy maker receives the information in form of probability distributions over demand functions and outcomes.

4.1 Consumer

At each point in time, the representative agent consumes quantities $c_{i,t}$ of a continuum of differentiated products $i \in [0,1]$ and provides labor L_t to firms. We assume that the information about demand can be parameterized via a utility representation. It is important to note that by this we do not assume the availability of information about the agent's utility scale. Any monotone transformation of the utility function would yield the same policy recommendation by the mutual information criterion. Therefore, to the extent that the predicted demand functions of the model agree with the predicted demand functions observed in the economy, the model will give correct policy recommendations.[7]

According to the information of the policy maker, the agent's preferences have the following utility representation:

$$C_t = C(\mathbf{c}_t) = \left(\int_i \xi_{i;t} (c_{i;t})^{\sigma} \right)^{\frac{1}{\sigma}}, \quad 0 < \sigma < 1; \tag{4}$$

$$U_t = U(C_t, L_t) = \frac{C_t^{1-\theta}}{1-\theta} - \frac{(L_t)^{\phi+1}}{\phi+1}, \quad 0 < \theta < 1, \ \ \phi > 0; \tag{5}$$

$$V_t(\mathbf{U}) = \sum_{j=0}^{\infty} \beta^{t+j} U_{t+j} \tag{6}$$

For each individual good, the policy maker's uncertainty about the consumer's preference is represented by a stochastic preference parameter $\xi_{i;t} \sim \ln N(\mu_\xi, \sigma_\xi^2)$, determining the agent's taste for the good in this time period.

The agent then maximizes $E_t V_t(\mathbf{U})$ by choosing a sequence[8] of choice vari-

[7] In fact, behavior only needs to be observationally indistinguishable from the predicted demand functions. Thus, a behavioral agent who has no well defined demand function but who's behavior is observationally indistinguishable from the policy maker's perspective can be treated as maximizing this demand function. This is guaranteed by the Observational Equivalence Axiom.

[8] Here and in the following, $\{\dots\}_{j=0}^{\infty}$ always denote sequences. We will omit subscripts whenever the index is implicitly understood.

ables $\{\mathbf{c}_{t+j}, L_{t+j}\}_{j=0}^{\infty}$, where $\mathbf{c}_t = \{c_{i;t}\}_{i\in[0,1]}$, subject to the budget constraint:

$$w_t L_t + B_{t-1}R_{t-1} + d_t = \int_i p_{i;t}c_{i;t} + B_t, \tag{7}$$

where w_t is the wage level in period t, B_t is the amount of bonds holding at the end of period t, R_t is the nominal gross return between period t and $t+1$, and d_t is the dividend received from firms.

As in standard DSGE models, optimality conditions of the consumption side can be summarized as follows:

— Good i's individual demand function:

$$c_{i;t} = \left(\frac{p_{i;t}}{P_t}\right)^{\frac{1}{\sigma-1}} (\xi_{i;t})^{\frac{1}{1-\sigma}} C_t \tag{8}$$

— The Euler equation:

$$(C_t)^{-\theta} = \beta E_t\left[\frac{P_t}{P_{t+1}} R_t (C_{t+1})^{-\theta}\right] \tag{9}$$

— The labor supply function:

$$\frac{L_t^{\phi}}{C_t^{-\theta}} = \frac{w_t}{P_t} \tag{10}$$

where $P_t \equiv \left(\int_i p_{i;t}^{\frac{\sigma}{\sigma-1}} \xi_{i;t}^{\frac{1}{1-\sigma}}\right)^{\frac{\sigma-1}{\sigma}}$ is the unit price of aggregate consumption C_t.

4.2 Firms

For the production sector, we assume that each individual firm i is a monopoly supplier with the following linear production function,

$$q_{i;t} = f(L_{i;t}; \delta_{i;t}, \Delta_t) = \delta_{i;t}\Delta_t(L_{i;t}), \tag{11}$$

where $q_{i;t}$ is firm i's output in period t. A firm's productivity is stochastic in the sense that it is subject to both a individual-wise shock $\delta_{i;t}$, and a economy-wise shock Δ_t in any period. We assume that both $\delta_{i;t}$ and Δ_t are log-normally distributed with $\delta_{i;t} \sim \ln N(\mu_\delta, \sigma_\delta^2)$ and $\Delta_t \sim \ln N(\mu_\delta, \sigma_\Delta^2)$. Each firm i demands $L_{i,t}$ units of labor from the total amount of labor supplied. All firms offer the same wage to the agent.

Our model employs Calvo (1983) price staggering; in every time period, each firm has a probability of $1-\alpha$ to adjust its price for good i and otherwise uses the previous period's price. Under this assumption, aggregate price dynamics are given by:

$$\begin{aligned} P_t &= \left[\int_i p_{i;t}^{\frac{\sigma}{\sigma-1}} \xi_{i;t}^{\frac{1}{1-\sigma}}\right]^{\frac{\sigma-1}{\sigma}} \\ &= \left[(1-\alpha)\int_i {p_{i;t}^*}^{\frac{\sigma}{\sigma-1}} \xi_{i;t}^{\frac{1}{1-\sigma}} + \alpha P_{t-1}^{\frac{\sigma}{\sigma-1}}\right]^{\frac{\sigma-1}{\sigma}}. \end{aligned} \tag{12}$$

Notice that in case $\alpha = 0$, we are in the classical model with flexible prices.

Firm i's maximization problem at time t when facing the opportunity to reset its price is:

$$\max_{p_{i;t}^*} E_t \sum_{j=0}^{\infty} \alpha^j Q_{t,t+j} m_{t+j}(p_{i;t}^*), \tag{13}$$

where

$$m_{t+j}(p_{i;t}^*) = p_{i;t}^* \left(\frac{p_{i;t}^*}{P_{t+j}}\right)^{\frac{1}{\sigma-1}} \xi_{i;t+j}^{\frac{1}{1-\sigma}} C_{t+j} - \frac{w_{t+j}}{\delta_{i;t}\Delta_t}\left(\frac{p_{i;t}^*}{P_{t+j}}\right)^{\frac{1}{\sigma-1}} \xi_{i;t+j}^{\frac{1}{1-\sigma}} C_{t+j} \tag{14}$$

denotes its profit in period $t+j$ given the current reset price, and $Q_{t,t+j} = \beta^j \left(\frac{C_{t+j}}{C_t}\right)^{-\theta}\left(\frac{P_t}{P_{t+j}}\right)$ is the stochastic discount factor.

The first order condition of (13) is given by

$$p_{i;t}^* = \frac{E_t \sum_{j=0}^{\infty} \alpha^j Q_{t,t+j} P_{t+j}^{\frac{1}{1-\sigma}} \xi_{i;t+j}^{\frac{1}{1-\sigma}} C_{t+j} \frac{w_{t+j}}{\delta_{i;t+j}\Delta_{t+j}}}{\sigma E_t \sum_{j=0}^{\infty} \alpha^j Q_{t,t+j} P_{t+j}^{\frac{1}{1-\sigma}} \xi_{i;t+j}^{\frac{1}{1-\sigma}} C_{t+j}} \tag{15}$$

$$= \frac{1}{\sigma} \frac{E_t \sum_{j=0}^{\infty} (\alpha\beta)^j \Phi_{t+j} \xi_{i;t+j}^{\frac{1}{1-\sigma}} \frac{w_{t+j}}{\delta_{i;t+j}\Delta_{t+j}}}{E_t \sum_{j=0}^{\infty} (\alpha\beta)^j \Phi_{t+j} \xi_{i;t+j}^{\frac{1}{1-\sigma}}}, \tag{16}$$

where $\Phi_{t+j} \equiv P_{t+j}^{\frac{1}{1-\sigma}} C_{t+j}^{1-\theta}$ denotes aggregate economic dynamics in period $t+j$.

Notice that unless $\alpha = 0$, $p_{i;t}$ cannot be solely expressed by exogenous variables. We instead rely on log-linearization for the determination of the optimal reset price. Recall that (8) tells us that $c_{i;t}$, the variable on which our freedom measure focuses, depends on both $p_{i;t}$ and P_t. Therefore, we focus on the linearization of the relative price $\frac{p_{i;t}}{P_t}$ around the steady state such that,

$$\hat{\delta}_{i;t} \equiv \ln \frac{\delta_{i;t}}{\bar{\delta}} \approx \frac{\delta_{i;t} - \bar{\delta}}{\bar{\delta}} = 0,$$
$$\hat{\Delta}_t \equiv \ln \frac{\Delta_t}{\bar{\Delta}} \approx \frac{\Delta_t - \bar{\Delta}}{\bar{\Delta}} = 0,$$
$$\hat{\xi}_{i;t} \equiv \ln \frac{\xi_{i;t}}{\bar{\xi}} \approx \frac{\xi_{i;t} - \bar{\xi}}{\bar{\xi}} = 0,$$

$\hat{w}_t = \bar{w}$, $p_{i;t+j} = \bar{p} = \frac{\bar{w}}{\sigma\bar{\delta}\bar{\Delta}}$, $P_{t+j} = \bar{P} = \bar{p}\bar{\xi}^{\frac{-1}{\sigma}}$, and $C_{t+j} = \bar{C}$. As a simplifying assumption, there is no price heterogeneity in the steady state.

Log-linearization of the optimal relative price $\frac{p_{i;t^*}}{P_t}$ gives:

$$\ln \frac{p^*_{i;t}}{P_t} \approx \frac{1}{\sigma} \ln \bar{\xi} + \sum_{j=1}^{\infty} E_t(\alpha\beta)^j \hat{\pi}_{t+j} \tag{17}$$
$$+ (1-\alpha\beta) \sum_{j=0}^{\infty} E_t(\alpha\beta)^j \left(\hat{w}_{t+j} - \hat{P}_{t+j} - \hat{\delta}_{i;t+j} - \hat{\Delta}_{t+j}\right), \tag{18}$$

where we define similarly that $\hat{\pi}_{t+j} \equiv \ln \frac{P_{t+j}}{P_{t+j-1}}$ and $\hat{P}_{t+j} = \ln \frac{P_{t+j}}{\bar{P}}$. Notice that since $\delta_{i;t}$ and Δ_t are i.i.d across time, we have $E_t\hat{\delta}_{i;t+j} = 0$, $E_t\hat{\Delta}_{t+j} = 0$ for $j > 0$ according to our definition of $\bar{\delta}$ and $\bar{\Delta}$. We may then simplify (18) as:

$$\ln \frac{p^*_{i;t}}{P_t} \approx \frac{1}{\sigma} \ln \bar{\xi} + (1-\alpha\beta)(-\hat{\delta}_{i;t} - \hat{\Delta}_t) \tag{19}$$
$$+ \sum_{j=0}^{\infty} E_t(\alpha\beta)^j \hat{\pi}_{t+j} + (1-\alpha\beta) \sum_{j=1}^{\infty} E_t(\alpha\beta)^j \left(\hat{w}_{t+j} - \hat{P}_{t+j}\right) \tag{20}$$

Next, we linearize the aggregate price dynamics (12) and get the following relationship between inflation and optimal relative price:

$$\hat{\pi}_t = \frac{1-\alpha}{\alpha} \ln(\bar{\xi})^{\frac{-1}{\sigma}} + \frac{1-\alpha}{\alpha} \int_i \ln \frac{p^*_{i;t}}{P_t} - \frac{(1-\alpha)}{\alpha\sigma} \int_i \hat{\xi}_{i;t} \tag{21}$$

This, together with (20), gives the following equation about the inflation and the productivity shock:

$$\hat{\pi}_t = \frac{1-\alpha}{\alpha}(1-\alpha\beta) \left(\hat{w}_t - \hat{P}_t - \hat{\Delta}_t - \int_i \hat{\delta}_{i;t}\right) + \beta E_t \hat{\pi}_{t+1} - \frac{1-\alpha}{\alpha} \frac{1-\alpha\beta}{\sigma} \int_i \hat{\xi}_{i;t} \tag{22}$$

We employ Uhlig (1996)'s law of large numbers, in which the integral over a continuum of i.i.d random variables can be treated as the mean of that random variable. That is, $\int_i \hat{\delta}_{i;t} = E_t(\hat{\delta}_{i;t}) = 0$ and $\int_i \hat{\xi}_{i;t} = E_t(\hat{\xi}_{i;t}) = 0$ according to our definition of the steady state. This guarantees that we can ignore the effects of individual shocks and good-specific shocks on the whole economy. We then simplify (22) as:

$$\hat{\pi}_t = \frac{1-\alpha}{\alpha}(1-\alpha\beta)\left(\hat{w}_t - \hat{P}_t - \hat{\Delta}_t\right) + \beta E_t \hat{\pi}_{t+1}, \tag{23}$$

which is simply the New-Keynesian Philips Curve in terms of real marginal cost. Finally we plug (23) into (20) and gets the following expression for optimal relative price:

$$\ln \frac{p^*_{i;t}}{P_t} = \frac{1}{\sigma}\ln \bar{\xi} + \frac{\alpha}{1-\alpha}\hat{\pi}_t - (1-\alpha\beta)\hat{\delta}_{i;t} \tag{24}$$

4.3 Equilibrium

Having determined optimal consumer and firm behavior, we derive equilibrium conditions in terms of log-linearization. We begin with the labor market. The labor demand is given by:

$$\hat{L}^d_t = \frac{1}{\sigma-1}\frac{1}{\sigma}\ln\bar{\xi} + \hat{C}_t - \hat{\Delta}_t - \int_i \hat{\delta}_{i;t} + \frac{1}{1-\sigma}\int_i \hat{\xi}_{i;t} + \frac{1-\alpha}{\sigma-1}\sum_{j=0}\alpha^j \int_i \ln\frac{p^*_{i;t-j}}{P_t} \tag{25}$$

$$= \hat{C}_t - \hat{\Delta}_t \tag{26}$$

and labor supply is:

$$\hat{L}^s_t = \phi^{-1}\left(\hat{w}_t - \hat{P}_t - \theta\hat{C}_t\right) \tag{27}$$

Equating (26) and (27), and substituting (23) for $\hat{w}_t - \hat{P}_t$ gives the following equation:

$$\hat{C}_t = \frac{1+\phi}{\phi+\theta}\hat{\Delta}_t + \frac{\alpha}{1-\alpha}\frac{1}{\phi+\theta}\frac{1}{1-\alpha\beta}\left(\hat{\pi}_t - \beta E_t\hat{\pi}_{t+1}\right) \tag{28}$$

which, again, is the New-Keynesian Phillips Curve in terms of output. Another key equation describing the equilibrium is the log-linearized Euler Equation:

$$\hat{R}_t = \theta E_t(\hat{C}_{t+1} - \hat{C}_t) + E_t\hat{\pi}_{t+1}. \tag{29}$$

Equations (28) and (29) are the non-policy block of this DSGE model. Once a monetary policy is fixed, (28) and (29) determine the equilibrium inflation rate $\hat{\pi}_t$, together with the output level $\hat{C}_t$ and optimal relative price $\ln\frac{p^*_{i;t}}{P_t}$ accordingly. These are the major components in the expression of $c_{i;t}$ and therefore are essential in the calculation of our welfare criterion.

5 Freedom of Choice Under Various Policy Rules

We combine the freedom measure proposed in Section 3 and the monetary model of Section 4 and evaluate consumption freedom under various monetary policy rules. Notice that the individual's choice of demand functions is parametrized by $\xi_{i;t}$, the policy maker's information over the individual's preference for each good. In other words, every set of parameters $\xi_{i;t}$ for all goods specifies a unique demand function. Moreover, after aggregation across goods, each parameter $\xi_{i;t}$ only affects good i and no other good j. Thus, freedom of choice is represented by the mutual information between $\xi_{i;t}$ and $c_{i;t}$. The demand function of individual goods (8) tells us that

$$c_{i;t}(p_{i;t}) = \left(\frac{p_{i;t}}{P_t}\right)^{\frac{1}{\sigma-1}} \xi_{i;t}^{\frac{1}{1-\sigma}} C_t, \tag{30}$$

and we can infer that the individual's behavior is distorted by fluctuations in $\frac{p_{i;t}}{P_t}$ and C_t. From the perspective of the policy maker, the distribution of $\frac{p_{i;t}}{P_t}$ is a mixture distribution under the Calvo pricing assumption. For any randomly picked good, its price duration—the length between current period and its last reset opportunity—is stochastic from an ex-ante perspective. This means that $p_{i;t}$ is a mixture distribution whose components are the distributions of the optimal reset prices of past periods. We denote S as the mixture's indicator variable, with the probability of the price duration s being realized being $(1-\alpha)\alpha^s$ for all $s \in \mathbb{N}$, and

$$p_{i;t} = \begin{cases} p^*_{i;t} & \text{if S} = 0 \\ p^*_{i;t-1} & \text{if S} = 1 \\ \vdots & \vdots \\ p^*_{i;t-s} & \text{if S} = s \\ \vdots & \vdots \end{cases}.$$

It follows immediately that $\frac{p_{i;t}}{P_t}$ as well as $c_{i;t}(p_{i;t})$ are also mixture distributions. As will be shown later, while each mixture component $c_{i;t}(p^*_{i;t-s})$ belongs to the same family of probability distributions, the mixture itself in general is not of any common probability distribution. This makes it difficult to obtain an analytical expression of the mutual information between $\xi_{i;t}$ and $c_{i;t}(p_{i;t})$ directly. Ideally, we would like to compute the mutual information between $\xi_{i;t}$ and $c_{i;t}(p^*_{i;t-j})$ in each state and sum them up linearly. However, by doing so we are ignoring the information loss incurred from mixing. More precisely,[9]

$$\sum_{s=0}(1-\alpha)\alpha^s I\left(\xi_{i;t}; c_{i;t}(p^*_{i;t-s})\right) = I\left(\xi_{i;t}; c_{i;t} \mid \mathrm{S}\right) = I\left(\xi_{i;t}; c_{i;t}\right) + I(\xi_{i;t}; \mathrm{S} \mid c_{i;t}). \tag{31}$$

One can see that an additional term $I(\xi_{i;t}; S \mid c_{i;t})$ needs to be subtracted from the LHS to obtain $I(\xi_{i;t}; c_{t;t})$. As we do not know the distribution of $c_{i;t}$, without specifying a particular monetary policy, this adjustment term cannot be obtained analytically. Fortunately, we can still make some qualitative remarks about it. Decomposing this adjustment term yields:

$$I(\xi_{i;t}; S \mid c_{i;t}) = H(S \mid c_{i;t}) - H(S \mid c_{i;t}, \xi_{i;t}). \tag{32}$$

The conditional entropy terms depend on the divergence between mixture components of $c_{i;t}$.[10] We consider the following two cases:

— Low divergence: $H(S \mid c_{i;t}, \xi_{i;t}) \approx H(S \mid c_{i;t})$. Low divergence means that all component distributions of the mixture are almost identical, and thus $\xi_{i;t}$ affects each component almost identically. That is, $\xi_{i;t}$ is almost independent of S conditional on $c_{i;t}$. It follows that $H(S \mid c_{i;t}, \xi_{i;t}) \approx H(S \mid c_{i;t})$.

— High divergence: Under high divergence S is no longer conditionally independent of $\xi_{i;t}$ given $c_{i;t}$. However, a high divergence leads to a small $H(S \mid c_{i;t})$. This is because $c_{i;t}$ becomes a perfect predictor of S if the supports of $c_{i;t-s}$ differ greatly for different values of s. It then follows that $I(\xi_{i;t}; S \mid c_{i;t}) = H(S \mid c_{i;t}) - H(S \mid c_{i;t}, \xi_{i;t}) \leq H(S \mid c_{i;t})$ is also very small.

The effect of the support of mixture components (the effect of parameters that affect macroeconomic stability) is not necessarily monotonic and may be nonzero for intermediate divergence. However, since the effect is zero in both extreme cases, we ignore this term as a secondary effect. Qualitatively, this means that we focus on the disturbances of the relation between demand and outcomes from price variation at the reset time and ignore variations due to different reset times.

[9] See appendix C.

[10] As shown later, $c_{i;t}(p^*_{i;t-s})$ have the same mean for all $s \geq 0$. Thus when speaking of their divergence, we refer to the differences between their variances.

5.1 Constant Price

The *divine coincidence* discussed by Blanchard and Galí (2007) states that monetary policies face no trade-off between price and output gap stabilization in a new Keynesian model without real rigidities. Here we revisit this neutrality with a different welfare criterion that does not only depend on inflation and the output gap. We begin with a price-stabilizing rule in which the inflation rate is kept zero.

Proposition 1. *Under a zero inflation policy rule $\hat{\pi}_t = 0$,*

$$I(\xi_{i;t}; c_{i;t})^{CP} = I\left(\xi_{i;t}; c_{i;t}(p^*_{i;t-s})\right)^{CP} = \frac{1}{2}\ln\left(1 + \frac{\left(\frac{1}{1-\sigma}\right)^2 \sigma_\xi^2}{\left(\frac{1-\alpha\beta}{1-\sigma}\right)^2 \sigma_\delta^2 + \left(\frac{1+\phi}{\phi+\theta}\right)^2 \sigma_\Delta^2}\right) \tag{33}$$

From the proposition we can see that mutual information is increasing in σ_ξ^2 and decreasing in σ_δ^2 and σ_Δ^2. First, freedom increases if the diversity of choices is higher. A low σ_ξ^2 means that the individual's preferences over different goods are almost identical. A high σ_ξ^2 means that the individual may prefer, for example, high food consumption and low recreational consumption, or low food consumption and high recreational consumption.

σ_δ^2 and σ_Δ^2 capture the degree of fluctuation in the production of goods. σ_δ^2 determines the volatility of the marginal rate of transformation between goods and σ_Δ^2 captures the fluctuation of overall output. The more stochastic the production conditions are, the less control the individual has over the outcomes. That is, if σ_δ^2 is high, the individual's consumption depends more on whether a particular firm is very productive or not; if σ_Δ^2 is high consumption depends on whether the whole economy is very productive or less productive.

Finally, we observe that freedom is increasing in both α and β. α determines the level of rigidity, that controls the extent to which production conditions impact prices. When the rigidity is high, relative prices automatically remain relatively stable and less severe interventions are necessary to achieve price stability. The previous effect may be further magnified by β, as it captures the patience of the individual, and more importantly, the firms. Under the Calvo pricing assumption, firms take profits in future periods into account when adjusting prices, and patient firms adjust their prices less aggressively to current production conditions. Once again, stable relative prices distort consumption less significantly.

5.2 Constant Output

Having evaluated the case with price-stabilization, we now examine a policy in which the aggregate output level is kept constant.

Proposition 2. *Under a policy rule that fully stabilizes output, $\hat{C}_t = 0$,*

$$I(\xi_{i;t}; c_{i;t}(p^*_{i;t-s}))^{CO} = \frac{1}{2}\ln\left(1 + \frac{\sigma_\xi^2}{(1-\alpha\beta)^2\sigma_\delta^2 + (1-\alpha\beta)^2(1+\phi)^2\left[1 + s\left(\frac{1-\alpha}{\alpha}\right)^2\right]\sigma_\Delta^2}\right), \tag{34}$$

and

$$I(\xi_{i;t}; c_{i;t})^{CO} \approx \sum_{s=0}^{\infty}(1-\alpha)\alpha^s I(\xi_{i;t}; c_{i;t}(p^*_{i;t-s}))^{CO} \tag{35}$$

As we discussed above, whenever we obtain a mixture distribution, overall freedom of choice can only be approximated. We analyze the comparative statics of the terms in (34) to determine the effect on the approximation in (35). In the constant output policy, each parameter acts quite similarly to the previous case with a few exceptions. First, θ becomes irrelevant since it controls the intertemporal substitutability of consumption; if output is stabilized, it has no effect. Second, compared to the constant price equilibrium, an increase in ϕ has the opposite effect. In the constant output equilibrium, labor supply decreases in $\hat{\Delta}_t$. In order to incentivize the individual to work enough hours when the economy is unproductive, the wage level increases. A low level of intertemporal labor substitution implies that the wage, as well as the aggregate price level, has to be driven even higher to achieve the same effect on output. This means that the effect of $\hat{\Delta}_t$ on inflation is increased for larger ϕ. From these differences, we can already foresee that the optimal choice between output and price stabilization will depend on parameters. We analyze this tradeoff in the next subsection. In the subsection thereafter, we generalize the analysis to allow for a continuum of policies between output and price stabilization.

5.3 Comparison of Price and Output Stabilization

The following corollary provides some insight into comparison between output and price stability as a monetary policy:

Corollary 2.1. $I(\xi_{i;t}; c_{i;t})^{CP} > I(\xi_{i;t}; c_{i;t})^{CO}$ *if* $\frac{1-\alpha\beta}{1-\sigma} > \frac{1}{\phi+\theta}$.

The choice between these two policies is equivalent to a trade-off between their drawbacks—volatile output versus volatile inflation. A change in a parameter makes a policy more desirable, if it alleviates the drawbacks induced by that policy. Corollary 2.1 tells us that a constant price policy induces more consumption freedom than a constant output policy when θ and ϕ are high. In contrast, when α and β are high or when σ is low, a constant price policy would become more desirable. The explanation goes as follows.

Firstly, higher θ and ϕ means that the individual keeps a smoother consumption pattern over time. When a constant price policy is implemented, the problem it raises —volatile output— is alleviated. Higher α and β in general make prices stabler. Thus, when a constant output policy is implemented, the problem of volatile inflation is less severe under high α and β. Lastly, a low σ means a lower degree of substitution between goods. This keeps $c_{i;t}$ closer to C_t; in other words, aggregate output has a stronger effect on the consumption of a particular good. If C_t is affected by Δ_t, as it is in price stabilization, this directly feeds through to $c_{i;t}$, leading to lower consumption freedom.

From the perspective of a utilitarian policy maker who knows the consumers' exact utility function, stabilizing total output is a mistake as the natural level of output is known. However, when agnostic about the utility function and only equipped with a model of behavior, the best the policy maker can do is to minimize external disturbance to the consumers' choices. For the policy maker it is irrelevant whether these disturbances originate from external shocks or from the monetary policy. If stabilizing prices leads to too large fluctuations in total output, the policy maker may prefer output stability over price stability.

5.4 Inflation targeting

The above two cases deal with rather extreme circumstances in which the policy maker coercively stabilizes one dimension of aggregate fluctuations. Empirically, such extreme policies are often unattainable. Instead, both the monetary policy literature and central banks in the real world consider interest rate rules preferable. We continue with the analysis of a zero-inflation targeting rule. Unlike the above two "constant" policies, an interest rate rule diverts the impact of aggregate shock into fluctuation of both price level and output.

Proposition 3. *Under an inflation targeting policy rule, $\hat{R}_t = \rho_\pi \hat{\pi}_t$, $\rho_\pi > 0$,*

$$I(\xi_{i;t}, c_{i;t})^{IT} = (1-\alpha) I\left(\xi_{i;t}, c_{i;t}(p^*_{i;t})\right)^{IT} + \sum_{s=1}^{\infty} (1-\alpha)\alpha^s I\left(\xi_{i;t}, c_{i;t}(p^*_{i;t-s})\right)^{IT}, \tag{36}$$

where

$$I\left(\xi_{i;t}, c_{i;t}(p^*_{i;t})\right)^{IT} = \frac{1}{2}\ln\left(1 + \frac{\sigma_\xi^2}{(1-\alpha\beta)^2\sigma_\delta^2 + (1-\sigma)^2\left(\frac{1+\phi}{\phi+\theta}\right)^2\left(\frac{\rho_\pi + \theta\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}{\rho_\pi+\theta A}\right)^2\sigma_\Delta^2}\right), \tag{37}$$

$$I\left(\xi_{i;t}, c_{i;t}(p^*_{i;t-s})\right)^{IT} = \frac{1}{2}\ln\left(1 + \frac{\sigma_\xi^2}{(1-\alpha\beta)^2\sigma_\delta^2 + \left(\frac{1+\phi}{\phi+\theta}\right)^2\left(\frac{\theta}{\rho_\pi+\theta A}\right)^2 B\sigma_\Delta^2}\right), \tag{38}$$

with $A = \frac{\alpha}{1-\alpha}\frac{1}{\phi+\theta}\frac{1}{1-\alpha\beta}$ *and* $B = \left[\left(\frac{(1-\sigma)\rho_\pi - \theta}{\theta}\right)^2 + \left(\frac{\alpha}{1-\alpha}\right)^2 + (s-1)\right]$.

As it turns out, the optimal policy can be derived using first order conditions with respect to ρ. We then obtain the following comparative statics of the optimal intensity ρ^* with which the policy maker should pursue price stability.

Corollary 3.1. *Let ρ^* denote the policy coefficient that maximize $I\left(\xi_{i;t}, c_{i;t}\right)^{IT}$. We have $\frac{\partial\rho^*}{\partial\phi} > 0$, $\frac{\partial\rho^*}{\partial\theta} > 0$, $\frac{\partial\rho^*}{\partial\sigma} > 0$, $\frac{\partial\rho^*}{\partial\beta} < 0$, and $\frac{\partial\rho^*}{\partial\alpha} < 0$,*

A high ρ represents a stronger emphasis on stabilizing inflation, which in equilibrium corresponds to lower fluctuations in prices and higher fluctuations in absolute output. One can easily observe that the effects of parameters coincide with those in Corollary 2.1, and the same explanation is valid here as well. The comparison between the constant output policy and constant price policy therefore yields the same qualitative results as the local comparative statics of the optimal interest rate rule. A greater emphasis should be placed on price stability if firms are impatient, the economy has low rigidities, goods are highly substitutable, and the intertemporal substitutability of consumption and leisure is low.

6 Conclusion

Two limitations of our study directly lead to possibilities for future research. In our study, we focused on presenting a parsimonous model that allows to gain insight into how our introduced welfare criterion responds to behavioral parameters and how maximizing this objective differs from maximizing a utilitarian criterion. It is well known that the model we employed does not fit behavioral patterns observed in the economy (for example, there is no inflation persistence). A natural extension is therefore to consider our welfare criterion in a larger model that fits observed behavior better. Moreover, there is an interesting dichotomy in the effects of inflation on individuals. Small fluctuations in inflation and output tend to primarily affect consumers, while large fluctuations tend to affect producers and their employees due to imperfect financial markets. Our model therefore is no longer valid once inflation and output fluctuations affect labor market choices via firm bankruptcies. A more detailed analysis of this would have led us further from the standard new Keynesian model. Using our welfare criterion on labor market/production choices in addition to consumption choices is therefore another natural extension of our analysis.

We conclude with a summary of the three main contributions of this paper. First, we provided a novel framework in which the question of the constitutional choice for a central bank can be analyzed without relying on cardinally comparable utility functions while still obtaining a complete ordering of the policies. Second, we showed that the policy maker may put emphasis on absolute output stabilization and that therefore the presence of the divine coincidence in the model does not imply that a zero inflation policy is optimal. Third, we obtained comparative statics for optimal monetary constitutions represented by output stabilization, price stabilization, and inflation targeting. We found that in an economy with low price rigidities, impatient firms, highly substitutable goods, and a low intertemporal substitutability a greater emphasis should be placed on price stability than output stability.

References

Akerlof, G. A. (2002). Behavioral Macroeconomics and Macroeconomic Behavior. *American Economic Review*, *92*(3), 23.

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

Blanchard, O., Dell'Ariccia, G., & Mauro, P. (2010, August 18). Rethinking Macroeconomic Policy. *Journal of Money, Credit and Banking*, *42*, 199–215. doi:10.1111/j.1538-4616.2010.00334.x

Blanchard, O. & Galí, J. (2007, January 18). Real Wage Rigidities and the New Keynesian Model. *Journal of Money, Credit and Banking*, *39*, 35–65. doi:10.1111/j.1538-4616.2007.00015.x

Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., & Jansen, D.-J. (2008, November). Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence. *Journal of Economic Literature*, *46*(4), 910–945. doi:10.1257/jel.46.4.910

Boettke, P. J. & Smith, D. J. (2016, December). Evolving views on monetary policy in the thought of Hayek, Friedman, and Buchanan. *The Review of Austrian Economics*, *29*(4), 351–370. doi:10.1007/s11138-015-0334-8

Calvo, G. A. (1983, September). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, *12*(3), 383–398. doi:10.1016/0304-3932(83)90060-0

Clarida, R., Gali, J., & Gertler, M. (1999). The Science of Monetary Policy: A New Keynesian Perspective. *Journal of Economic Literature*, *37*, 150.

Coenen, G. (2007, January). Inflation persistence and robust monetary policy design. *Journal of Economic Dynamics and Control*, *31*(1), 111–140. doi:10.1016/j.jedc.2005.09.012

D'Amico, D. (2007, November 9). Buchanan on monetary constitutions. *Constitutional Political Economy*, *18*(4), 301–318. doi:10.1007/s10602-007-9027-z

Debortoli, D., Kim, J., Lindé, J., & Nunes, R. (2019, July 1). Designing a Simple Loss Function for Central Banks: Does a Dual Mandate Make Sense? *The Economic Journal*, *129*(621), 2010–2038. doi:10.1111/ecoj.12630

Dowding, K. & van Hees, M. (2009). Freedom of Choice. In *Handbook of Rational and Social Choice*. Oxford, UK: Oxford University Press.

Edge, R. M., Laubach, T., & Williams, J. C. (2010, January). Welfare-maximizing monetary policy under parameter uncertainty. *Journal of Applied Econometrics*, *25*(1), 129–143. doi:10.1002/jae.1136

Elster, J. & Roemer, J. E. (1991). *Interpersonal Comparisons of Well-Being*. Cambridge University Press.

Flood, R. P. & Isard, P. (1989). Monetary Policy Strategies. *IMF Staff Papers*, *36*(3), 612–632.

Friedman, B. M. (1988). Lessons on Monetary Policy from the 1980s. *Journal of Economic Perspectives*, *2*(3), 51–72.

Giannoni, M. P. (2002, February). Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model. *Macroeconomic Dynamics*, *6*(1), 111–144. doi:10.1017/S1365100502027062

Goodfriend, M. (2007). How the World Achieved Consensus on Monetary Policy. *Journal of Economic Perspectives*, *21*(4), 47–68.

Grant, S., Kajii, A., Polak, B., & Safra, Z. (2010). Generalized Utilitarianism and Harsanyi's Impartial Observer Theorem. *Econometrica*, *78*(6), 1939–1971. doi:10.3982/ECTA6712

Harsanyi, J. C. (1953). Welfare Economics of Variable Tastes. *The Review of Economic Studies*, *21*(3), 204. doi:10.2307/2295773

Harsanyi, J. C. (1955). Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility. *Journal of Political Economy*, *61*, 434–435.

Harsanyi, J. C. (1977). *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge, UK: Cambridge University Press.

Kraskov, A., Stögbauer, H., & Grassberger, P. (2004). Estimating mutual information. *Physical review E*, *69*(6), 66138.

Kydland, F. E. & Prescott, E. C. (1977, June). Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, *85*(3), 473–491. doi:10.1086/260580

Levin, A. T. & Williams, J. C. (2003, July). Robust monetary policy with competing reference models. *Journal of Monetary Economics*, *50*(5), 945–975. doi:10.1016/S0304-3932(03)00059-X

Lohmann, S. (1992). Optimal Commitment in Monetary Policy: Credibility versus Flexibility. *American Economic Review*, *82*(1), 273–286.

Orphanides, A. (2003a, April). Monetary policy evaluation with noisy information. *Journal of Monetary Economics*, *50*(3), 605–631. doi:10.1016/S0304-3932(03)00027-8

Orphanides, A. (2003b, April). The quest for prosperity without inflation. *Journal of Monetary Economics*, *50*(3), 633–663. doi:10.1016/S0304-3932(03)00028-X

Orphanides, A. & Williams, J. C. (C. (2002). Robust Monetary Policy Rules with Unknown Natural Rates. *Brookings Papers on Economic Activity*, *2002*(2), 63–145. doi:10.1353/eca.2003.0007

Rommeswinkel, H. (2019). Measuring Freedom in Games.

Rotemberg, J. & Woodford, M. (1997). Optimization-based Econometric Framework for the Evaluation of Monetary Policy. *NBER Macroeconomics annual*, *12*(1), 297–346.

Rotemberg, J. & Woodford, M. (1999). Interest rate rules in an estimated sticky price model. In *Monetary policy rules* (pp. 57–126). Chicago University Press. Retrieved from http://www.nber.org/chapters/c7414.pdf

Rötheli, T. F. (2015). Elements of behavioral monetary economics. In *Handbook of contemporary behavioral economics* (pp. 711–727). Routledge.

Sala, L., Söderström, U., & Trigari, A. (2008, July). Monetary policy under uncertainty in an estimated model with labor market frictions. *Journal of Monetary Economics*, *55*(5), 983–1006. doi:10.1016/j.jmoneco.2008.03.006

Sen, A. (1977). Non-linear social welfare functions: A reply to Professor Harsanyi. In *Foundational problems in the special sciences* (pp. 297–302). Springer.

Shiller, R. J. (2003). From efficient markets theory to behavioral finance. *Journal of economic perspectives*, *17*(1), 83–104.

Sims, C. A. (2003, April). Implications of rational inattention. *Journal of Monetary Economics*, *50*(3), 665–690. doi:10.1016/S0304-3932(03)00029-1

Taylor, J. B. (Ed.). (1999). *Monetary Policy Rules*. Chicago: University of Chicago Press, v. 31. Studies in Business Cycles.

Uhlig, H. (1996). A law of large numbers for large economies. *Economic Theory*, 10.

Walsh, C. E. (1995). Optimal Contracts for Central Bankers. *American Economic Review*, *85*(1), 19.

Weymark, J. A. (2011). A reconsideration of the Harsanyi-Sen debate on utilitarianism. In *Interpersonal Comparisons of Well-Being*. Cambridge, UK: Cambridge University Press.

Appendices

A Formal Definitions

In this section, we give formal definitions of all the axioms and concepts introduced in the main text.

We denote the set of demand functions of the consumer by $\mathcal{X} = \{(p, w) \mapsto (c, y) : pc = yw\}$. We assume that there exists a set of outcomes $\mathcal{O}$ that consist of the quantities $c \in \mathbb{R}^I_+$ of the goods indexed by I consumed by the agent and the hours worked by the agent, $y \in \mathbb{R}$. The set of policies, $\mathcal{P}$ can be assumed to be equal to the set $\mathcal{P}$ of finite support probability measures on $(\mathcal{X} \times \mathcal{O})$ endowed with the product topology. The preferences of the policy maker are a binary relation $\succsim \subseteq \mathcal{P} \times \mathcal{P}$.

Expected Utility Rationality

Axiom A.1 (Rationality)**.** For all $P, P', P'' \in \mathcal{P}$,

— $P \succsim P'$ or $P' \succsim P$, or both.

— If $P \succsim P'$ and $P' \succsim P''$, then $P \succsim P''$.

Axiom A.2 (Continuity)**.** For any two convergent sequences of policies, $\{P_k\}_{k=1}^{\infty} \to P$ and $\{P'_k\}_{k=1}^{\infty} \to P'$, if for all k, $P_k \succsim P'_k$, then $P \succsim P'$.

To define expected utility rationality of the policy maker, we define lotteries as follows:

Definition 1 (Lotteries)**.** The set of lotteries is $\mathcal{P}_L = \{P \in \mathcal{P} : \forall x \in \mathcal{X}, o \in \mathcal{O} : P[o, x] = P[o]P[x]\}$.

Axiom A.3 (Lottery Independence)**.** If $P, P', P'' \in \mathcal{P}_L$, then $P \succsim P'$ if and only if $\alpha P \oplus (1-\alpha)P'' \succsim \alpha P \oplus (1-\alpha)P''$ for all $\alpha \in (0, 1)$.

Observational Equivalence:

Definition 2 (Observationally Equivalent Demand)**.** Two demand functions x, x' are observationally equivalent under policy P, denoted $x \approx_P x'$, if $P[o|x] = P[o|x']$ for all $o \in \mathcal{O}$.

We can therefore denote the partition of $\mathcal{X}$ into sets of observationally equivalent demand functions as $\mathcal{X}/\approx_P$.

Definition 3 (Observationally Equivalent Policies)**.** Two policies P, P' are observationally equivalent, denoted $P \approx P'$ if for all sets $\bar{\mathcal{X}} \in \mathcal{X}/\approx_P$, there exists a set $\bar{\mathcal{X}}' \in \mathcal{X}/\approx_{P'}$ such that $P[\mathcal{X}] = P[\mathcal{X}']$ and for $x \in \bar{\mathcal{X}}$ and $x \in \bar{\mathcal{X}}'$, $P[o|x] = P'[o|x']$ for all $o \in \mathcal{O}$.

Axiom A.4 (Observational Equivalence)**.** $P \approx P'$ implies $P \sim P'$ for all $P, P' \in \mathcal{P}$.

Choice Deprivation:

Definition 4 (Choice Deprivation)**.** $P' = D^{\mathcal{X}'}P$ is the result of a choice deprivation if

$$P'[x,o] = (D^{\mathcal{X}'}P)[x,o] = \begin{cases} P[x]\sum_{x' \in \mathcal{X}'} P[x',o], & x \in \mathcal{X}' \\ P[x,o], & \text{else} \end{cases} \tag{39}$$

Alternatively, choice deprivations can be seen as fulfilling the following conditions.

— $P[o,x] = P'[o,x]$ for all $x \in \mathcal{X} \backslash \mathcal{X}'$ and all $o \in \mathcal{O}$.

— $P[o|\mathcal{X}'] = P'[o|\mathcal{X}']$ for all $o \in \mathcal{O}$.

— $P[x] = P'[x]$ for all $x \in \mathcal{X}'$.

Based on this definition, we define independence of the preference in identical choice deprivations as follows.

Axiom A.5 (Deprivation Independence)**.** For any two policies $P, P' \in \mathcal{P}$, if $P[x,o] = P'[x,o]$ for all $x \in \mathcal{X}' \subset \mathcal{X}$ and all $o \in \mathcal{O}$, then,

$$P \succsim P' \quad \Leftrightarrow \quad D^{\mathcal{X}'}P \succsim D^{\mathcal{X}'}P'. \tag{40}$$

Decomposability: Formally, we define a subpolicy as the policy obtained by Bayesian updating on a set of outcomes $\mathcal{O}'$:

$$P[x,o|\mathcal{O}'] = \begin{cases} \frac{P[x,o]}{P[\mathcal{O}']} & o \in \mathcal{O}' \\ 0, & \text{else.} \end{cases} \tag{41}$$

We introduce the following formal definition of decomposability.

Axiom A.6 (Weak Decomposability)**.** If $P[\cdot|\mathcal{O} - \mathcal{O}'] \approx P'[\cdot|\mathcal{O} - \mathcal{O}']$, and $P[\mathcal{O}'] = P[\mathcal{O}']$, then: $P \succsim P'$ if and only if $P[\cdot, \mathcal{O}'] \succsim P'[\cdot, \mathcal{O}']$.

Mathematically, this states that common subpolicies can be ignored when comparing two policies, only the subpolicies that are distinct are relevant for the preference between the policies.

B Proof of Theorem 1

Proof. We use the result of Rommeswinkel (2019), Theorem 1, to prove this result. We first obtain a representation of each policy P as a process (G_P, θ_P) that Rommeswinkel (2019) defines as combinations of game forms with probability distributions over strategies. Thus, we map the set of policies into a larger space of combinations of game forms with finite support probability measures over mixed strategies on the game form. The cited result then provides a representation theorem under which we obtain the desired representation. We therefore need to show that the axioms imposed on the set of policies imply the axioms of Rommeswinkel (2019) on the set of processes. For this, we need to show that there is a mapping from policies into the processes, f. From the preference $\succsim$ on the set of policies we can then define the preference $\succsim^*$ as $f[P] \succsim^* f[P']$ if $P \succsim P'$. Under this mapping, the axioms on $\succsim$ in this paper must then translate into the axioms on $\succsim^*$ in Rommeswinkel (2019).

For every policy P, we define a game form G_P as follows. $G_P = (\{1\}, \mathcal{X}, o_P)$ where $\mathcal{X}$ is the set of demand functions. o_P is a mapping from demand functions into lotteries over outcomes $\mathcal{O}$ defined by:

$$(o_P)[c, l] = P[(c, l)|x] \tag{42}$$

In other words, the lottery resolved in the game form G_P after the agent has chosen demand function x is the conditional probability distribution of the outcomes given the demand function derived from the policy. Finally, we assume the following distribution θ over mixed strategies:

$$\theta_P[\mathbb{1}_x] = P[x] \tag{43}$$

where $\mathbb{1}_x$ is the probability measure that yields x with certainty. In other words, in the process the probability measure θ over strategies is such that the player plays a pure strategy. We now define the mapping f as $f[P] = (G_P, \theta_P)$.

The mapping $f : P \mapsto G_P$ is not surjective, since the set of processes contains processes with mixed strategies. However, it is one-to-one, since any distinction in P and P' results either in distinct o_P and $o_{P'}$ or distinct θ_P and $\theta_{P'}$. However, under the Outcome Equivalence axiom of Rommeswinkel (2019), every process with mixed strategies is indifferent to a process with pure strategies in which each pure strategy yields the same conditional probability of outcomes via the lottery o_P. Thus, we can uniquely extend $\succsim^*$ to the set of all processes such that the Outcome Equivalence axiom is maintained.

Having obtained a unique preference $\succsim^*$, we can verify that it fulfills the axioms Rationality, Continuity, Lottery Independence, Outcome Equivalence, Strategy Independence, and Subprocess Monotonicity.

Rationality is straightforward. After extending $\succsim^*$ to processes with mixed strategies, the relation is complete. Moreover, transitivity of $\succsim$ directly translates into transitivity of $\succsim^*$.

Continuity of $\succsim^*$ requires that the weakly lower $\{\theta' : (G, \theta) \succsim^* (G, \theta')\}$ and upper sets $\{\theta' : (G, \theta') \succsim^* (G, \theta)\}$ of the relation are closed. The Continuity axiom on $\succsim$ guarantees that the sets $\{P' : P \succsim P'\}$ and $\{P' : P' \succsim P\}$ are closed (all sequences in the upper and lower sets converge and the space is metric). Given that all $\{P' : P \succsim P'\}$ and $\{P' : P' \succsim P\}$ are closed and $P \mapsto o_P$ and $P \mapsto \theta_P$ are continuous, it follows that all $\{\theta' : (G, \theta) \succsim^* (G, \theta')\}$ and upper sets $\{\theta' : (G, \theta') \succsim^* (G, \theta)\}$ are indeed closed.

Outcome Equivalence of processes with mixed strategies holds by definition. For all remaining processes, we use Observational Equivalence. Outcome Equivalence requires that all processes are indifferent in which the equivalence classes of strategies with the same conditional probabilities of outcomes are equally likely. For two processes (G, θ) and (G', θ'), it follows that for any equivalence class $\mathcal{X}' \subseteq \mathcal{X}$, $\theta[\mathcal{X}'] = \theta'[\mathcal{X}']$. But then by the definition of f, $P[\mathcal{X}'] = P'[\mathcal{X}']$ where $f(P) = (G, \theta)$ and $f(P') = (G', \theta')$. Observational equivalence guarantees that whenever for two policies P, P', if their equivalence classes of demand functions (that yield the same conditional probabilitiy of outcomes) are equally likely, then they are indifferent. It follows from the definition of $\succsim^*$ that then also (G, θ) and (G', θ') are indifferent, proving Outcome Equivalence.

The Lottery Independence axiom on $\succsim^*$ follows from Lottery Independence on $\succsim$ and the fact that policies that are lotteries translate via f into processes that contain no influential player.

Strategy Independence follows directly from Deprivation Independence after realizing that $D_1^{\mathcal{M}_1} f[P] = f[D^{\mathcal{X}'} P]$ where $\mathcal{M}_1 = \{\mathbb{1}_x : x \in \mathcal{X}'\}$.

Subprocess Monotonicity on $\succsim^*$ is implied by Weak Decomposability of $\succsim$ and Outcome Equivalence, which has above been shown to hold. Under Outcome Equivalence, for every process (G, θ) we can find an outcome equivalent process $g[G, \theta]$ such that each action yields only one particular outcome with certainty. A subprocess is obtained by conditioning the strategies to the subgame and conditioning θ on the subgame. In processes in which every action yields a unique outcome, every set of actions is a disjoint subgame, since there is only a single player. Therefore, conditioning on subgames is equivalent to conditioning on sets of outcomes. Therefore, the conditioning on a subprocess to a subgame containing outcomes $\mathcal{O}'$ is the same as conditioning a policy to the set $\mathcal{O}'$. If P' is a subpolicy of P obtained by conditioning on the set of outcomes $\mathcal{O}'$, then $g[f[P']]$ is a subprocess of $g[f[P]]$ obtained by conditioning on the subgame of the actions that each yield one of the outcomes $\mathcal{O}'$ with certainty. Since Weak Decomposability requires that $\succsim$ is monotone in the ranking of a subpolicy, it follows that $\succsim^*$ is monotone in subprocesses.

We have therefore shown that the axioms imposed on $\succsim$ imply the axioms imposed on $\succsim^*$. By Theorem 1 of Rommeswinkel (2019), it follows that $\succsim^*$ has a representation of the form:

$$U[f[P]] = \sum_{x\in\mathcal{X}} \theta_P[x] \sum_{c\in\mathcal{O}} (o_P[x])[c] \left(v[c] + r \cdot \ln \frac{(o_P[x])[c]}{\sum_{x'} \theta_P[x'](o_P[x'])[c]} \right) \tag{44}$$

and thus $\succsim$ has the desired representation. □

C Derivation of (31)

The sum of mutual information between $\xi_{i;t}$ and the mixture components $\{c^*_{i;t-j}\}_{j\in\mathbb{N}}$ can be expressed as:

$$\begin{aligned}
\sum_{j\in\mathbb{N}} (1-\alpha)\alpha^j I(\xi_{i;t}; c_{i;t}(p^*_{i;t-j})) = &\sum_{j\in\mathbb{N}} (1-\alpha)\alpha^j H(\xi_{i;t}) \\
&- \sum_{j\in\mathbb{N}} (1-\alpha)\alpha^j H(\xi_{i;t} \mid c_{i;t}(p^*_{i;t-j})) && (45)\\
=&H(\xi_{i;t}) - \sum_{j\in\mathbb{N}} (1-\alpha)\alpha^j H(\xi_{i;t} \mid c_{i;t}(p^*_{i;t-j})) && (46)\\
=&H(\xi_{i;t}) - \sum_{j\in\mathbb{N}} P(J=j) H(\xi_{i;t} \mid c_{i;t}, J=j) && (47)\\
=&\underbrace{H(\xi_{i;t} \mid J)}_{\because \xi_{i;t} \perp\!\!\!\perp J} - H(\xi_{i;t} \mid c_{i;t}, J) && (48)\\
=&I(\xi_{i;t}; c_{i;t} \mid J), && (49)
\end{aligned}$$

which is actually the mutual information between $\xi_{i;t}$ and the mixture conditional on price duration. Next, by the chain rule of mutual information,

$$\begin{aligned}
I(\xi_{i;t}; c_{i;t}, J) &= I(\xi_{i;t}; J) + I(\xi_{i;t}; c_{i;t} \mid J) && (50)\\
&= I(\xi_{i;t}; c_{i;t}) + I(\xi_{i;t}; J \mid c_{i;t}). && (51)
\end{aligned}$$

Since $I(\xi_{i;t}; J) = 0$, we have,

$$I(\xi_{i;t}; c_{i;t}) = I(\xi_{i;t}; c_{i;t} \mid J) - I(\xi_{i;t}; J \mid c_{i;t}) \tag{52}$$
$$= \sum_{j \in \mathbb{N}} (1-\alpha)\alpha^j I(\xi_{i;t}; c_{i;t}(p^*_{i;t-j})) - I(\xi_{i;t}; J \mid c_{i;t}). \tag{53}$$

D Proof of Proposition 1

Proof. The equilibrium under policy rule $\hat{\pi}_t = 0$, and the non-policy block conditions (28), (29) can be summarized as:

$$\hat{C}_t = \frac{1+\phi}{\phi+\theta}\hat{\Delta}_t \tag{54}$$
$$\ln \frac{p^*_{i;t-s}}{P_{t-s}} = \frac{1}{\sigma}\ln \bar{\xi} - (1-\alpha\beta)\hat{\delta}_{i;t-s}, \tag{55}$$

which can be plugged into (8) to obtain the log-demand function $\ln c_{i;t}(p_{i;t})$, whose mixture components are:

$$\ln c_{i;t}(p^*_{i;t-s}) = \frac{1}{\sigma-1}\left(\ln \frac{p^*_{i;t-s}}{P_{t-s}} - \sum_{k=1}^{s} \hat{\pi}_{t-s+k}\right) + \frac{1}{1-\sigma}\ln \xi_{i;t} + \ln C_t \tag{56}$$
$$= \frac{1}{\sigma-1}\left(\frac{1}{\sigma}\ln \bar{\xi} - (1-\alpha\beta)\hat{\delta}_{i;t-s}\right) + \frac{1}{1-\sigma}\ln \xi_{i;t} + \left(\frac{1+\phi}{\phi+\theta}\hat{\Delta}_t + \bar{C}\right). \tag{57}$$

Note that each mixture component $\ln c_{i;t}(p^*_{i;t-s})$ is a combination of normally distributed variables, and thus $\ln \xi_{i;t}$ and $\ln c_{i;t}(p^*_{i;t-s})$ should be jointly normal. We denote their covariance matrix as:

$$\Sigma_{t-s} = \begin{pmatrix} \sigma^2_{\xi} & \\ \frac{\sigma^2_{\xi}}{1-\sigma} & V + \left(\frac{1+\phi}{\phi+\theta}\right)^2 \sigma^2_{\Delta} \end{pmatrix}, \tag{58}$$

where $V = \left(\frac{1-\alpha\beta}{1-\sigma}\right)^2 \sigma^2_{\delta} + \left(\frac{1}{1-\sigma}\right)^2 \sigma^2_{\xi}$.

The mutual information between two jointly normal random variables N, N' is given by:

$$I(N, N') = \frac{1}{2}\ln\left(\frac{Var[N]Var[N']}{|\Sigma_{N,N'}|}\right). \tag{59}$$

Additionally, as suggested by Kraskov, Stogbauer, & Grassberger (2004), mutual information is invariant under homeomorphism. Thus,

$$
\begin{aligned}
I(\xi_{i;t}; c_{i;t}) &\approx \sum_{s=0} (1-\alpha)\alpha^s I\left(\xi_{i;t}; c_{i;t}(p^*_{i;t-s})\right) && (60)\\
&= \sum_{s=0} (1-\alpha)\alpha^s I\left(\ln \xi_{i;t}; \ln c_{i;t}(p^*_{i;t-s})\right) && (61)\\
&= \frac{1}{2}\ln\left(\frac{\sigma_\xi^2\left(V+\left(\frac{1+\phi}{\phi+\theta}\right)^2\sigma_\Delta^2\right)}{\sigma_\xi^2\left(V+\left(\frac{1+\phi}{\phi+\theta}\right)^2\sigma_\Delta^2\right)-\left(\frac{\sigma_\xi^2}{1-\sigma}\right)^2}\right) && (62)\\
&= \frac{1}{2}\ln\left(1+\frac{\left(\frac{1}{1-\sigma}\right)^2\sigma_\xi^2}{\left(\frac{1-\alpha\beta}{1-\sigma}\right)^2\sigma_\delta^2+\left(\frac{1+\phi}{\phi+\theta}\right)^2\sigma_\Delta^2}\right). && (63)
\end{aligned}
$$

□

E Proof of Proposition 2

Proof. First, plug the policy rule $\hat{C}_t = 0$ into (28) and solve the difference equation for $\hat{\pi}_t$, with

$$\hat{\pi}_t = -\frac{1-\alpha}{\alpha}(1-\alpha\beta)(1+\phi)\hat{\Delta}_t. \tag{64}$$

It then follows that from (24),

$$\ln\frac{p^*_{i;t-s}}{P_{t-s}} = \frac{1}{\sigma}\ln\bar{\xi} - (1-\alpha\beta)(1+\phi)\hat{\Delta}_{t-s} - (1-\alpha\beta)\hat{\delta}_{i;t-s}. \tag{65}$$

Mixture components of the log-demand function then takes the following form:

$$
\begin{aligned}
\ln c_{i;t}(p^*_{i;t-j}) = \frac{1}{\sigma-1}&\left[\frac{1}{\sigma}\ln\bar{\xi} - (1-\alpha\beta)(1+\phi)\hat{\Delta}_{t-j} - (1-\alpha\beta)\hat{\delta}_{t-j}\right. && (66)\\
&\left.+\sum_{s=1}^{j}\frac{1-\alpha}{\alpha}(1-\alpha\beta)(1+\phi)\hat{\Delta}_{t-s+k}\right] + \frac{1}{1-\sigma}\ln\xi_{i;t} + \ln\bar{C}, && (67)
\end{aligned}
$$

whose covariance matrix with $\ln\xi_{i;t}$ is:

$$\Sigma_{t-s} = \begin{pmatrix} \sigma_\xi^2 & \\ \frac{\sigma_\xi^2}{1-\sigma} & V + (1+\phi)^2\left(\frac{1-\alpha\beta}{1-\sigma}\right)^2\left[1+s\left(\frac{1-\alpha}{\alpha}\right)^2\right]\sigma_\Delta^2 \end{pmatrix}, \tag{68}$$

where $V = \left(\frac{1-\alpha\beta}{1-\sigma}\right)^2 \sigma_\delta^2 + \left(\frac{1}{1-\sigma}\right)^2 \sigma_\xi^2$. The mutual information between $\xi_{i;t}$ and $c_{i;t}(p^*_{i;t-s})$ is then

$$MI(c_{i;t}(p^*_{i;t-s}); \xi_{i;t}) = \frac{1}{2} \ln \left(\frac{\sigma_\xi^2 \left[V + (1+\phi)^2 \left(\frac{1-\alpha\beta}{1-\sigma}\right)^2 \left[1 + s\left(\frac{1-\alpha}{\alpha}\right)^2 \right] \sigma_\Delta^2 \right]}{\sigma_\xi^2 \left[V + (1+\phi)^2 \left(\frac{1-\alpha\beta}{1-\sigma}\right)^2 \left[1 + s\left(\frac{1-\alpha}{\alpha}\right)^2 \right] \sigma_\Delta^2 \right] - \left(\frac{\sigma_\xi^2}{1-\sigma}\right)^2} \right) \tag{69}$$

$$= \frac{1}{2} \ln \left(1 + \frac{\sigma_\xi^2}{(1-\alpha\beta)^2 \sigma_\delta^2 + (1-\alpha\beta)^2 (1+\phi)^2 \left[1 + s\left(\frac{1-\alpha}{\alpha}\right)^2 \right] \sigma_\Delta^2} \right). \tag{70}$$

□

F Proof of Proposition 3

Proof. Plugging the policy rule $\hat{R}_t = \rho_\pi \hat{\pi}_t$ into (28) and (29) gives

$$\hat{\pi}_t = \frac{-\theta}{\rho_\pi + \theta A} \frac{\phi+1}{\phi+\theta} \hat{\Delta}_t, \tag{71}$$

$$\hat{C}_t = \frac{\phi+1}{\phi+\theta} \left(\frac{\rho_\pi}{\rho_\pi + \theta A} \right) \hat{\Delta}_t, \tag{72}$$

where $A = \frac{\alpha}{1-\alpha} \frac{1}{\phi+\theta} \frac{1}{1-\alpha\beta}$.
It then follows that by (24),

$$\ln \frac{p^*_{i;t-s}}{P_{t-s}} = \ln \bar{\xi}^{\frac{1}{\sigma}} - \frac{\alpha}{1-\alpha} \frac{\theta}{\rho_\pi + \theta A} \frac{\phi+1}{\phi+\theta} \hat{\Delta}_{t-s} - (1-\alpha\beta)\hat{\delta}_{i;t-s} \tag{73}$$

Mixture components of the log-demand function then takes the following form:

$$\begin{aligned} \ln c_{i;t}(p^*_{i;t-s}) = \frac{1}{\sigma-1} \Bigg(& \ln \bar{\xi}^{\frac{1}{\sigma}} - \frac{\alpha}{1-\alpha} \frac{\theta}{\rho_\pi + \theta A} \frac{\phi+1}{\phi+\theta} \hat{\Delta}_{t-s} - (1-\alpha\beta)\hat{\delta}_{i;t-s} \\ & + \frac{\theta}{\rho_\pi + \theta A} \frac{\phi+1}{\phi+\theta} \sum_{k=1}^{j} \hat{\Delta}_{t-j+k} \Bigg) + \frac{1}{1-\sigma} \ln \xi_{i;t} + \frac{\phi+1}{\phi+\theta} \left(\frac{\rho_\pi}{\rho_\pi + \theta A} \right) \hat{\Delta}_t + \bar{C}, \end{aligned} \tag{74}$$

whose covariance matrices with $\ln \xi_{i;t}$ are

$$\Sigma_t = \begin{pmatrix} \sigma_\xi^2 & \\ \frac{\sigma_\xi^2}{1-\sigma} & V + \left(\frac{\phi+1}{\phi+\theta}\right)^2 \left(\frac{1}{\rho_\pi+\theta A}\right)^2 \left(\rho_\pi + \theta \frac{\alpha}{1-\alpha} \frac{1}{1-\sigma}\right)^2 \sigma_\Delta^2 \end{pmatrix} \tag{75}$$

for flexible-price goods ($s = 0$) and

$$\Sigma_{t-s}\big|_{j\neq 0} = \begin{pmatrix} \sigma_\xi^2 \\ \frac{\sigma_\xi^2}{1-\sigma} & V + \left(\frac{\phi+1}{\phi+\theta}\right)^2 \left(\frac{1}{\rho_\pi+\theta A}\right)^2 \left[\left(\rho_\pi + \frac{\theta}{\sigma-1}\right)^2 + \left(\frac{\alpha}{1-\alpha}\frac{\theta}{1-\sigma}\right)^2 + (s-1)\left(\frac{\theta}{1-\sigma}\right)^2\right]\sigma_\Delta^2 \end{pmatrix} \tag{76}$$

for non-flexible-price goods ($s > 0$) respectively. The mutual information between $\xi_{i;t}$ and $c_{i;t}(p^*_{i;t-s})$ is then:

$$I\left(\xi_{i;t}, c_{i;t}(p^*_{i;t})\right)^{IT} = \frac{1}{2}\ln\left(1 + \frac{\sigma_\xi^2}{(1-\alpha\beta)^2\sigma_\delta^2 + (1-\sigma)^2\left(\frac{1+\phi}{\phi+\theta}\right)^2\left(\frac{\rho_\pi+\theta\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}{\rho_\pi+\theta A}\right)^2\sigma_\Delta^2}\right), \tag{77}$$

$$I\left(\xi_{i;t}, c_{i;t}(p^*_{i;t-s})\right)^{IT} = \frac{1}{2}\ln\left(1 + \frac{\sigma_\xi^2}{(1-\alpha\beta)^2\sigma_\delta^2 + \left(\frac{1+\phi}{\phi+\theta}\right)^2\left(\frac{\theta}{\rho_\pi+\theta A}\right)^2 B\sigma_\Delta^2}\right), \tag{78}$$

with $A = \frac{\alpha}{1-\alpha}\frac{1}{\phi+\theta}\frac{1}{1-\alpha\beta}$ and $B = \left[\left(\frac{(1-\sigma)\rho_\pi-\theta}{\theta}\right)^2 + \left(\frac{\alpha}{1-\alpha}\right)^2 + (s-1)\right]$.

□

G Proof of Corollary 3.1

Proof. We begin the proof with several lemmas that are straightforward to derive.

Lemma 1. $\frac{\partial I\left(\xi_{i;t};c_{i;t}(p^*_{i;t})\right)^{IT}}{\partial\rho_\pi} \gtreqless 0$ *if* $A \lesseqgtr \frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}$.

Proof. We note that (77) is decreasing in the second summation term in the denominator of the fraction. Taking the derivative and simplifying directly yields the result. □

Lemma 2. $\forall s > 0$, $\rho^*_s = \frac{\theta}{(1-\sigma)A+1}\left[A + \frac{1}{1-\sigma}\left(s + \frac{\alpha}{1-\alpha}\right)^2\right]$ *maximizes* $I\left(\xi_{i;t}; c_{i;t}(p^*_{i;t-s})\right)^{IT}$.

Proof. We note that (78) is decreasing in the denominator of the fraction. Taking first order conditions and verifying concavity of the objective yields the desired solution. □

Lemma 3. $\frac{\partial\rho^*_s}{\partial\phi} > 0$, $\frac{\partial\rho^*_s}{\partial\theta} > 0$, $\frac{\partial\rho^*_s}{\partial\sigma} > 0$, $\frac{\partial\rho^*_s}{\partial\alpha} < 0$, $\frac{\partial\rho^*_s}{\partial\beta} < 0$, *and* $\frac{\partial\rho^*_s}{\partial s} > 0$.

Proof. These results follow directly from the solution of ρ^*_s. □

Lemma 4. *Let $\rho^*_{s,A\lesseqgtr\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$ denote the policy coefficient that maximizes $I\left(\xi_{i;t}, c_{i;t}(p^*_{i;t-s})\right)^{IT}$ when $A \lesseqgtr \frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}$. We have $\rho^*_{s,A<\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}} > \rho^*_{s,A>\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$.*

Proof. Rearranging $A > \frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}$ gives $\frac{1-\sigma}{(\phi+\theta)(1-\alpha\beta)} > 1$. We can see that the direction of the inequality changes when ϕ, θ, σ increases and α, β decreases, *ceteris peribus*. $\rho^*_{s,A<\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$ should then be greater than $\rho^*_{s,A>\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$ according to lemma 3. □

Lemma 5. *Let $\rho^*_{A\lesseqgtr\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$ denote the policy coefficient that maximizes $I\left(\xi_{i;t}, c_{i;t}\right)^{IT}$ when $A \lesseqgtr \frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}$. We have $\rho^*_{A<\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}} > \rho^*_{A>\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$.*

Proof. For simplicity, we hereafter denote $I\left(\xi_{i;t}, c_{i;t}(p^*_{i;t-s})\right)^{IT}$ as $I(*)_s$. Since $I\left(\xi_{i;t}, c_{i;t}\right)^{IT}$ does not have an analytical expression and therefore we cannot directly obtain its partial derivative, we first decompose $I\left(\xi_{i;t}, c_{i;t}\right)^{IT}$ as:

$$I(\xi_{i;t}, c_{i;t})^{IT} = (1-\alpha)I(*)_0 + \sum_{s=1}^{\infty}(1-\alpha)\alpha^s I(*)_s \tag{79}$$

$$\equiv (1-\alpha)I(*)_0 + I(*)_{1,\infty}. \tag{80}$$

Lemma 6. *Let $\rho^*_{1,\infty,A\lesseqgtr\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$ be the policy coefficient that maximizes $I(*)_{1,\infty}$ when $A \lesseqgtr \frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}$, then $\rho^*_{1,\infty,A<\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}} > \rho^*_{1,\infty,A>\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$.*

Proof. By lemma 4 we know that $\rho^*_{s,A<\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}} > \rho^*_{s,A>\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$ $\forall s > 0$. Since $I(*)_{1,\infty}$ is a additive function of $I(*)_s$'s, we have $\rho^*_{1,\infty,A<\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}} > \rho^*_{1,\infty,A>\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$. □

According to lemma 1, $I(*)_0$ cannot be maximized if $A < \frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}$; we then know from (80) that $\rho^*_{A<\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}} > \rho^*_{1,\infty,A<\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$. On the other hand, if $A > \frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}$, lemma 1 tells us that $I(*)_0$ is maximized when $\rho = 0$; thus (80) tells that $\rho^*_{A>\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}} < \rho^*_{1,\infty,A>\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$. According to Claim 1, we then have:

$$\rho^*_{A<\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}} > \rho^*_{1,\infty,A<\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}} > \rho^*_{1,\infty,A>\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}} > \rho^*_{A>\frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}}$$

□

By the above lemma we know that the policy maker should choose a higher policy coefficient if $A < \frac{\alpha}{1-\alpha}\frac{1}{1-\sigma}$, namely if $D(\phi, \theta, \sigma, \alpha, \beta) \equiv \frac{1-\sigma}{(\phi+\theta)(1-\alpha\beta)} < 1$. To infer how each parameter affects the policy coefficient ρ^*, it is then sufficient to show how each parameter affects $D(\dots)$ so that the inequality should be satisfied. We can easily see that $\frac{\partial D}{\partial \phi} < 0$, $\frac{\partial D}{\partial \theta} < 0$, $\frac{\partial D}{\partial \sigma} < 0$, $\frac{\partial D}{\partial \beta} > 0$, and $\frac{\partial D}{\partial \alpha} > 0$, which means that higher ϕ, θ, and σ lead to a higher policy coefficient, and higher β and α lead to a lower policy coefficient.

$\square$