

Distributional Comparative Statics

Martin Kaae Jensen^{*†}

October 2016

Abstract

Distributional comparative statics is the study of how individual decisions and equilibrium outcomes vary with changes in the distribution of economic parameters (income, wealth, productivity, information, etc.). This paper develops new tools to address such issues and illustrates their usefulness in applications. The central development is a condition called *quasi-concave differences* which implies concavity of the policy function in optimization problems without imposing differentiability or quasi-concavity conditions. The general take-away is that many distributional questions in economics which cannot be solved by direct calculations or the implicit function theorem, can be addressed easily with this paper's methods. Several applications demonstrate this: the paper shows how increased uncertainty affects the set of equilibria in Bayesian games; it shows how increased dispersion of productivities affects output in the model of Melitz (2003); and it generalizes Carroll and Kimball (1996)'s result on concave consumption functions to the Aiyagari (1994) setting with borrowing constraints.

Keywords: Distributional comparative statics, concave policy functions, income distribution, inequality, uncertainty, heterogenous firms, Bayesian games, dynamic stochastic general equilibrium models, arg max correspondence.

JEL Classification Codes: C61, D80, D90, E20, I30.

*Department of Economics, University of Leicester. (e-mail: mj182@le.ac.uk)

†I would like to thank the managing editor Marco Ottaviani, four anonymous referees, Jean-Pierre Drugeon, Charles Rahal, Alex Rigos, Kevin Reffett, Colin Rowat, Jaideep Roy, John Quah, Muhamet Yildiz, the participants in my 2015 PhD course on comparative statics at the Paris School of Economics, and especially Daron Acemoglu and Chris Wallace for suggestions and comments that have influenced this paper substantially. Also thanks to participants at the 2012 European Workshop on General Equilibrium Theory in Exeter and seminar participants at Arizona State University, Humboldt University of Berlin, Paris School of Economics, University of Leicester, University of Zurich, and University of Warwick. All remaining errors are my responsibility.

1 Introduction

Distributional comparative statics (DCS) studies how changes in exogenous distributions affect endogenous distributions in models with optimizing agents. Examples of DCS questions are the following:

- Consider a monetary policy committee (MPC) that sets the interest rate by minimizing a standard loss function, as in Kydland and Prescott (1977). The public knows the MPC's objective but the MPC has private information about how the interest rate affects output and inflation. How does the accuracy of the public's beliefs about the MPC's private signal (the exogenous distribution) affect the public's interest rate expectations (the endogenous distribution)?
- The incomplete markets model of Aiyagari (1994) features a population of consumers with heterogeneous incomes who make consumption and savings decisions subject to borrowing constraints. At any moment in time, labor incomes differ due to inequality in endowments/productivities and the distribution of these is exogenous. Under what conditions on consumers' preferences does a Lorenz decrease in inequality reduce welfare according to a concave and increasing social welfare criterion? Under what conditions will a decrease in inequality raise per-capita savings?
- In the international trade model of Melitz (2003) a continuum of firms have different productivities. Does increased dispersion in productivity increase or reduce total output? *I.e.*, are there increasing or decreasing returns to diversity?
- In Bayesian games, the obvious DCS question is how less precise private signals affect the Bayesian equilibria, including the actions' means and variability. In an arms race with incomplete information, countries may be uncertain about arms' effectiveness and opponents' intentions — and the degree of uncertainty is likely to change over time. Does increased uncertainty lead to disarmament or to escalation in the Bayesian equilibrium? In more general Bayesian games, how does the precision of private signals affect mean and variability of Bayesian equilibrium actions?

As explained in Section 2 and illustrated repeatedly throughout this paper, concavity or convexity of the functions that map exogenous variables into endogenous ones (the policy functions) is the key to answering such questions whenever the distributions change in the sense of mean-preserving spreads, second-order stochastic, Lorenz, generalized Lorenz dominance, among other ways.¹ We may therefore focus our attention on concavity/convexity of these policy functions *provided* we know under what conditions on the primitives of the model they will be concave or convex. The main contribution of this paper is a theorem that shows that if a payoff function satisfies

¹This observation dates back to Atkinson (1970) who shows that in standard models of savings behavior, dominating shifts in Lorenz curves reduce or increase aggregate savings according to whether the savings function is concave or convex.

a condition called *quasi-concave differences*, then the policy function — and more generally, the policy correspondence — will be concave. The condition implies concavity of the policy function whether or not payoff functions are differentiable, concave, or even quasi-concave. This advances the literature in several ways.

First, the result enables us to deal with distributional issues in a number of models which we previously could not handle. Thus in the model of Aiyagari (1994) mentioned above, any attempt at using the implicit function theorem fails because the value function is not smooth (Section 2.2). In the trade setting of Melitz (2003), existing methods fail when production sets are not convex (Section 3.3), and so on.

Second, the isolation of the critical condition for DCS allows us to disentangle the fundamental economic conditions from unnecessary technical conditions. As this paper's applications illustrate again and again, this can improve our economic understanding substantially. Readers familiar with monotone methods (*e.g.* Topkis (1978), Milgrom and Shannon (1994), Quah (2007)) and with so-called robust comparative statics more generally (*e.g.* Milgrom and Roberts (1994), Acemoglu and Jensen (2015)) will immediately spot the parallel: when one obtains a result under certain sufficient conditions and those conditions are a mixture of critical economic conditions and unnecessary technical conditions, economic intuition is lost because one is unable to separate the two (Milgrom and Roberts (1994), p.442-443).² In contrast, this paper's results allow us to deal with changing distributions in full generality and see precisely under which conditions a specific conclusion holds.

Third, quasi-concave differences is easy to verify in applications. When functional forms are differentiable, it can be characterized explicitly via derivatives and is especially easy to work with.³ The alternative is "brute force", *i.e.*, to repeatedly apply the implicit function theorem on the first-order/Euler conditions. Often this leads to extremely cumbersome calculations. Whether brute force works or not, the tools developed here are usually much simpler to apply ("brute force" is discussed in detail in Section 2).

The paper begins in Section 2 by further motivating and exemplifying the DCS agenda. Section 2 also previews the paper's main results without going into too much technical detail. The paper then turns to quasi-concave differences, discusses the intuitive content of the definition, and shows — first in the simplest possible setting (Section 3.1), then under more general conditions (Section 3.2)— that quasi-concave differences implies concavity of the policy function in an optimization problem. An appendix treats the issue under yet more general conditions where the decision vector is allowed to live in an arbitrary topological vector lattice (Appendix C). Section 3.3 contains a practitioner's guide to the results and a fully worked-through example. Section 4 then tackles DCS in Bayesian games, and Section 5 derives general conditions for concavity of policy functions in stochastic dynamic programming problems. As a concrete application, Section 5 ex-

²Monotone methods have an important role to play in DCS (see Section 2.1), but they are rarely sufficient on their own. In particular, one cannot simply parameterize the exogenous distribution and then apply monotone methods (or the implicit function theorem).

³Operationally, the conditions for quasi-concave differences in the differentiable case are on an equal footing with, say, concavity, or supermodularity/increasing differences which can be established, respectively, via the Hessian criterion and the cross-partial derivatives test of Topkis (1978).

tends Carroll and Kimball (1996) to the setting with borrowing constraints (Aiyagari (1994)). That result plays an important role for various distributional comparative statics questions in macroeconomics (Huggett (2004), Acemoglu and Jensen (2015)) and is also essential for analyzing inequality in settings where consumers may be credit constrained.

2 Preview and Motivation

This section previews the paper’s results and explains the role of convex and concave policy functions for distributional comparative statics (DCS). The section also discusses several set-ups in which existing methods are unable to address DCS questions.

2.1 Forecasting Monetary Policy

A monetary policy committee (MPC) meets to set the rate of interest $x \in X \subseteq \mathbb{R}$. As in Kydland and Prescott (1977), the MPC has a loss function $L(y - y^*, \pi - \pi^*)$ where y denotes output, π denotes inflation, and stars denote natural/target levels. The central bank controls output and inflation via the interest rate, $y = y(x, z)$ and $\pi = \pi(x, z)$ where $z \in Z \subseteq \mathbb{R}$ is a parameter that represents the MPC’s private information about the Lucas supply curve.⁴ z may be thought of as the MPC’s estimation of future inflation, or it could reflect private information about the specific functional form of the Lucas supply curve. The MPC’s objective is thus to maximize $u(x, z) = -L(y(x, z) - y^*, \pi(x, z) - \pi^*)$ with respect to x .

A forecaster must predict the MPC’s interest rate decision. She knows its objective u but only has beliefs about z represented by a probability measure μ on Z . This implies a forecast with distribution

$$(1) \quad \mu_x(A) = \mu\{z \in Z : g(z) \in A\},$$

where A is any Borel set in X and $g : Z \rightarrow X$ is the MPC’s policy function,

$$(2) \quad g(z) = \arg \max_{x \in X} u(x, z).^5$$

So if the forecaster is asked how likely the MPC is to set the interest rate in the interval between 0.5 and 0.6 %, she will answer “with probability $\mu_x([0.5, 0.6])$ ” where $\mu_x([0.5, 0.6]) \in [0, 1]$. Her “headline” forecast will be the mean of μ_x . And so on.

Consider now a shift in the forecaster’s beliefs μ . For example, she might become more uncertain about the MPC’s private signal (a mean-preserving spread to μ) because the MPC decides to reveal less of its private information to the public. The interesting economic question is then how

⁴ A *Lucas supply curve* (LSC) describes the positive relationship between output y and inflation π for a given natural level of output and inflation expectations (see e.g. Heijdra (2009), Chapter 9). When the MPC controls y and π via the interest rate, it thus moves the economy up and down the Lucas supply curve. Note that we could have equally assumed that the central bank directly chooses inflation and output/unemployment as in Kydland and Prescott (1977). We would then calculate the interest rate that would accomplish these goals by means of the LSC. The reason for the focus on interest rates will become clear when the forecaster is introduced next.

⁵We assume here that the MPC is able to agree on a single decision (existence and uniqueness).

the distribution of the forecast μ_x changes in response. And most importantly, whether the headline forecast which investors and home buyers base their decisions on will increase or decrease. This paper's main result (Theorem 1) immediately allows us to conclude that if u satisfies a condition called *quasi-concave differences* (Definition 2 in Section 3), then the headline forecast will *decrease*. Quasi-concave differences says, roughly, that for any $\delta > 0$ close to zero, the difference between $u(x, z)$ and $u(x - \delta, z)$, $u(x, z) - u(x - \delta, z)$, must be quasi-concave. If $-u$ satisfies this, u exhibits *quasi-convex differences* and the headline forecast will instead increase. And this is true whether or not the objective is concave or even quasi-concave. The conditions are particularly easy to check when u is differentiable in x since then a key Lemma (Lemma 1) tells us that u exhibits quasi-concave differences if and only if the partial derivative $D_x u(x, z)$ is quasi-concave. If we specialize to an expected utility objective (think of \bar{z} as core inflation which is public information and has distribution η),

$$(3) \quad u(x, z) = \int U(x, \bar{z}, z) \eta(\bar{z}),$$

then it is easy to see that a sufficient condition for $D_x u(x, z)$ to be quasi-concave is that $D_x U(x, \bar{z}, z)$ is concave in (x, z) for almost every $\bar{z} \in Z$.⁶ Economically, concavity of $D_x U$ is equivalent to convexity of the MPC's marginal loss function $D_x L$. A convex marginal loss function obtains if the marginal loss is relatively constant or rises slowly when output and inflation are close to their target levels, and rises more rapidly when output and inflation are farther away from the targets.⁷ So if the MPC's adversity to an additional rate hike increases at an ever stronger rate the farther the MPC is from its targets, we should expect less information transmission to reduce mean forecasts.⁸

Now, increased uncertainty (mean preserving spreads) is just one type of belief shift that is of economic interest. The following Definition collects all of the stochastic orders considered in this paper (for an in-debt treatment see *e.g.* Shaked and Shanthikumar (2007)).

Definition 1 (Stochastic Orders) Let μ and $\tilde{\mu}$ be two distributions on the same measurable space $(Z, \mathcal{B}(Z))$.⁹ Also, let $f : Z \rightarrow \mathbb{R}$ be a function for which the following expression is well-defined,

$$(4) \quad \int f(z) \tilde{\mu}(dz) \geq \int f(z) \mu(dz).$$

Then:

- $\tilde{\mu}$ *first-order stochastically dominates* μ if (4) holds for any increasing function f .
- $\tilde{\mu}$ is a *mean-preserving spread* of μ if (4) holds for any convex function f .
- $\tilde{\mu}$ is a *mean-preserving contraction* of μ if (4) holds for any concave function f .¹⁰

⁶The details of everything being postulated here can be found in Section 3.1.

⁷Note that, strictly speaking, this interpretation requires that the Lucas supply curve is linear in x and z (*i.e.*, the functions $y(x, z)$ and $\pi(x, z)$ are linear). With non-linear relationships, the interest rate pass-through enters the picture and complicates matters. The topic of interpretation will occupy a large part of Section 3.1.

⁸For the related literature on central bank communication see *e.g.* Myatt and Wallace (2014) and references therein.

⁹Here $\mathcal{B}(Z)$ denotes the Borel algebra of Z .

¹⁰Note that $\tilde{\mu}$ is a mean-preserving contraction of μ if and only if μ is a mean-preserving spread of $\tilde{\mu}$.

- $\tilde{\mu}$ second-order stochastically dominates μ if (4) holds for any concave and increasing function f .
- $\tilde{\mu}$ dominates μ in the convex-increasing order if (4) holds for any convex and increasing function f .

Interesting economic examples abound in each case. For example, a second-order stochastic dominance increase could be due to an external event such as a more favorable public forecast of inflation or a drop in the price of oil. The relationship between such belief shifts and the forecast's distribution is simply a matter of inserting (1)-(2) into (4) and verify the condition of each bullet point in Definition 4. In each case, one immediately finds that the policy function g determines the outcome. This leads to the following Observations.¹¹ Note that by an "increase in μ ", we mean that μ is replaced with a distribution $\tilde{\mu}$ that dominates μ in the given stochastic order. Similarly for a "decrease in μ " and a "mean preserving spread to μ " where $\tilde{\mu}$ is dominated by μ and $\tilde{\mu}$ is a mean preserving spread of μ , respectively.

Observation 1. *If g is increasing, any first-order stochastic dominance increase in μ will lead to a first-order stochastic dominance increase in μ_x .*

Observation 2. *If g is concave, any mean-preserving spread to μ will lead to a second-order stochastic dominance decrease in μ_x .*

Observation 3. *If g is concave and increasing, any second-order stochastic dominance increase in μ will lead to a second-order stochastic dominance increase in μ_x .*

Observation 4. *If g is convex, any mean-preserving spread to μ will lead to a convex-increasing order increase in μ_x .*

Observation 5. *If g is convex and increasing, any convex-increasing order increase in μ will lead to a convex-increasing order increase in μ_x .*

Observation 1 tells us that if the MPC's policy function g is increasing, then a first-order stochastic dominance increase in the forecaster's beliefs μ implies a first-order stochastic dominance increase in the forecast's distribution μ_x . To show that g is increasing, one applies the implicit function theorem (IFT) or monotone methods. By the IFT, this is the case if $g'(z) \geq 0$ in equation (6) below. Using instead monotone methods, the same conclusion follows if u exhibits increasing differences (Topkis (1978)) or satisfies the single-crossing property (Milgrom and Shannon (1994)). So existing results fully enable us to deal with first-order stochastic dominance. And to be sure, the literature is full of instances of this argument.¹²

¹¹To be sure, there is nothing deep or difficult about these observations. They are as mentioned basically just restatements of the definitions. Nonetheless, detailed proofs are provided in Appendix A.

¹²For example, the property that first-order stochastic dominance of beliefs implies first-order stochastic dominance of (predicted) actions is the basic criteria for a Bayesian game to exhibit strategic complementarities, and Van Zandt and Vives (2007) provide multiple examples where they use monotone methods to verify that policy functions are increasing.

When we consider mean-preserving spreads (Observation 2) or second-order stochastic dominance (Observation 3), it is seen that g being increasing is not enough anymore. We must know whether g is concave to derive the effect on μ_x . Moreover, for the cases covered by Observations 4-5 we must know whether g is convex. This paper's results thus allow us to deal also with Observations 2-4. To the best of the author's knowledge, the only alternative is repeated use of the implicit function theorem (IFT). It is instructive to compare that approach to this paper's.

If U is sufficiently smooth, concavity of u is assumed, differentiation under the integral sign is allowed, and the solution is interior for all $z \in Z$, the following first-order condition is necessary and sufficient for an optimum

$$(5) \quad (D_x u(x, z) =) \int_{\bar{z} \in \bar{Z}} D_x U(x, \bar{z}, z) \eta(\bar{z}) = 0.$$

If the second derivative never equals zero (strict concavity of $u(\cdot, z)$), the IFT determines x as a function of z , $x = g(z)$ where

$$(6) \quad g'(z) = - \left[\int_{\bar{z} \in \bar{Z}} D_{xx}^2 U(g(z), \bar{z}, z) \eta(\bar{z}) \right]^{-1} \int_{\bar{z} \in \bar{Z}} D_{xz}^2 U(g(z), \bar{z}, z) \eta(\bar{z}).$$

Note that monotone comparative statics is about the sign of g' , and as Milgrom and Shannon (1994) convincingly argue, the IFT approach is not ideal for many applications. When the question is concavity of g , the situation is worse since we must determine g'' and so need to apply the IFT one more time. Specifically, we differentiate the right-hand-side of (6) with respect to z and substitute in for $g'(z)$. The resulting expression is rather daunting and of no particular importance to us. It contains a mixture of integrals of second and third derivatives and in contrast to the condition we arrived at using this paper's results above, it may or may not be possible to establish any useful and intuitively transparent condition for $g'' \leq 0$ (concavity) from such an expression.¹³ More substantially, in order to apply the IFT twice, a host of unnecessary technical assumptions must be imposed — so even when the IFT provides sufficient conditions for concavity of g , these will not be the most general conditions. As Milgrom and Roberts (1994) and Acemoglu and Jensen (2013, 2015) discuss in detail, this lack of “robustness” generally makes it impossible to disentangle the fundamental economic conditions that drive one's results from superfluous technical assumptions (again see also Milgrom and Shannon (1994), keeping in mind that the situation is worse here because we need to apply the IFT twice). In the current application, use of the IFT requires for example that the MPC's objective is strictly concave — which imposes spurious cross-restrictions on the loss function and the Lucas supply curve. If u is *not* strictly concave, or if it is *not* at least thrice differentiable the IFT is never applicable. We now turn to a particularly egregious instance of this.

2.2 Income Allocation Models

In the monetary policy committee example of the previous subsection, the implicit function theorem (IFT) does at least provide a conclusion under suitable technical assumptions. We now turn

¹³Note that the problem in part is that concavity of u simultaneously imposes conditions on second partial derivatives which leads to “entanglement” as discussed in the Introduction and further discussed momentarily. With multi-dimensional decision variables as explored in Appendix C, the IFT becomes excessively complicated and is rarely useful.

to an application from macroeconomics where the differentiability requirements of the IFT confounds *any* attempt to use it to establish concavity of the policy function. So here this paper's results provide the only known way to deal with the economic issues raised.

Consider a stochastic income allocation model with value function v and Bellman equation

$$(7) \quad v(x, z) = \max_{y \in \Gamma(x, z)} \tilde{u}((1+r)x + wz - y) + \beta \int v(y, z') \eta(dz').$$

$\Gamma(x, z) = \{y \in \mathbb{R} : -b \leq y \leq (1+r)x + wz\}$ is admissible savings given past savings x and labor productivity z which follows an *i.i.d.* process with distribution η . As usual r denotes the interest rate, and w the wage rate. The formulation where labor productivity is random follows Aiyagari (1994), but the discussion below — and indeed all of this paper's results — apply equally to the cases where the interest rate r is random or where both labor income and the interest rate are random (Carroll and Kimball (1996)). When $b < +\infty$, we have a borrowing constraint and potential market incompleteness (Aiyagari (1994), see also Acemoglu and Jensen (2015)).

Let $g((1+r)x + wz) = \arg \max_{y \in \Gamma(x, z)} \tilde{u}((1+r)x + wz - y) + \beta \int v(y, z') \eta(dz')$ denote the *savings function*, and $c((1+r)x + wz) = (1+r)x + wz - g((1+r)x + wz)$ the *consumption function*. These are, without further elaboration, assumed to be well-defined. Clearly, the savings function is convex if and only if the consumption function is concave.

Following Carroll and Kimball (1996), say that \tilde{u} belongs to the Hyperbolic Absolute Risk Aversion (HARA) class if $(\tilde{u}''' \cdot \tilde{u}') / (\tilde{u}'')^2 = k$ for a constant $k \in \mathbb{R}$. Carroll and Kimball (1996) prove that if \tilde{u} belongs to the HARA class, then the consumption function is concave if there is no borrowing constraint ($b = +\infty$) and if the period utility function has a positive third derivative (precautionary savings). The proof of Carroll and Kimball (1996) relies on Euler equations and repeated application of the IFT. In particular, it is necessary that the value function v is at least thrice differentiable. This is unproblematic if the borrowing constraint is inactive, but if $b < +\infty$, then the value function will not be thrice differentiable at any point where the borrowing constraint binds (Huggett (2004), p.773).¹⁴

As we shall see in Section 5, this paper's results allow for a simple and direct proof of the concavity of the consumption function which does not rely on Euler equations and does not require differentiability of the value function. In particular, it is shown that the consumption function will be concave for the general HARA class with or without borrowing constraints.¹⁵ Furthermore, the HARA class “pops out” endogenously from an application of our general results — there is no guesswork involved, and no ingenuity is required (contrast with the mathematical ingenuity of Carroll and Kimball (1996)). Note that this added simplicity when it comes to finding suitable sufficient conditions parallels the discussion of sufficient conditions for $g''(z) \leq 0$ that followed equation (6) in the previous subsection. In fact, the additional ease-of-use makes this paper's results so effective in the stochastic dynamic programming setting that little effort is required to

¹⁴See also Carroll and Kimball (2001) who address the concavity question in a framework with borrowing constraints in two special cases of the general HARA class (CRRA and CARA utility). Another contribution worth mentioning is Suen (2015) whose approach is entirely different and very powerful (the current paper and Suen (2015) were written independently of each other).

¹⁵As explained on page 25, the type of boundary condition imposed in the current paper does not rule out binding borrowing constraints.

prove a result on the convexity/concavity of policy functions for stochastic dynamic programming problems at the level of generality of the text book treatment of Stokey and Lucas (1989). Thus we are able in Section 5 to address not just the previous income allocation problem but nearly any stochastic dynamic model one can think of applying in macroeconomics and other fields.

3 Concave Policy Functions

Motivated by the previous section, we now present the paper's main results on the concavity and convexity of policy functions. The first subsection considers the simplest case of an objective with a one-dimensional decision variable, a fixed constraint set, and a unique optimizer. This simplicity allows us to focus on the new concepts' economic interpretation. In the second subsection, all of these restrictions are relaxed. The last subsection contains a user's guide to the results.

3.1 A Simple Case

Let $u : X \times Z \rightarrow \mathbb{R}$ be a payoff function where $x \in X \subseteq \mathbb{R}$ is a decision variable and $z \in Z$ a vector of parameters. It is assumed that X is convex and that Z is a convex subset of a real vector space. Assume also that the associated decision problem $\max_{x \in X} u(x, z)$ has a unique solution for all $z \in Z$.¹⁶ We may then define the *policy function* $g : Z \rightarrow X$

$$(8) \quad g(z) = \arg \max_{x \in X} u(x, z).$$

The model of Section 2.1 fits into this framework with $g(z)$ being the MPC's interest rate decision given its private information z . The purpose this section is to show that the policy function will be concave if the following condition holds.

Definition 2 (Quasi-Concave Differences) A function $u : X \times Z \rightarrow \mathbb{R}$ exhibits *quasi-concave differences* if for all $\delta > 0$ in a neighborhood of 0, $u(x, z) - u(x - \delta, z)$ is quasi-concave in $(x, z) \in \{x \in X : x - \delta \in X\} \times Z$.

If $-u$ exhibits quasi-concave differences, u exhibits *quasi-convex differences*. Quasi-convex differences will be shown to imply that g is convex. Note the close relationship between quasi-concave differences and Topkis (1978)'s notion of increasing differences.¹⁷ Quasi-concave differences is particularly easy to verify for differentiable objectives.

Lemma 1 (Differentiability Criterion) Assume that $u : X \times T \rightarrow \mathbb{R}$ is differentiable in $x \in X \subseteq \mathbb{R}$. Then u exhibits quasi-concave differences if and only if the partial derivative $D_x u(x, z)$ is quasi-concave in $(x, z) \in X \times Z$.

¹⁶For example, we have such a policy function if X is compact and u is upper semi-continuous and strictly quasi-concave in x . The relationship between strict quasi-concavity and this Section's main condition (quasi-concave differences) is discussed at the end of the section.

¹⁷In the current real vector space set-up, u exhibits increasing differences in x and z if and only if $u(x, z) - u(x - \delta, z)$ is (coordinatewise) increasing in z for all $x \in X$ and $\delta > 0$ with $x - \delta \in X$ (Topkis (1978)).

Proof. Appendix B. ■

As an illustration, consider the MPC's expected utility objective (3) from Section 2.1 with u continuously differentiable. We then have,

$$D_x u(x, z) = \int_{\bar{z} \in \bar{Z}} D_x U(x, \bar{z}, z) \eta(\bar{z}).$$

Since integration preserves concavity, it immediately follows from Lemma 1 that u exhibits quasi-concave differences if $D_x U(x, \bar{z}, z)$ is concave in (x, z) for *a.e.* \bar{z} .¹⁸

What is the economic interpretation of quasi-concave differences? To answer this question, consider a simple two-period, non-stochastic version of the income allocation model of Section 2.2. A consumer obtains utility $\tilde{u}(x, y)$ from consumption today $x \geq 0$ and consumption tomorrow $y \geq 0$. She receives income z today which she can either consume straight away or save for consumption tomorrow. The budget constraint is consequently $x + \frac{1}{1+r}y \leq z$ where r is the rate of interest she earns on her savings. When \tilde{u} is monotonically increasing in both arguments, and a boundary condition is imposed so that constraints can be ignored, this situation reduces to a one-dimensional decision problem with the structure of (8) by taking $u(x, z) = \tilde{u}(x, (1+r)(z-x))$ and writing $g(z) = \arg \max_{x \geq 0} u(x, z)$. In particular, $X = \mathbb{R}_+$ and $g(z)$ is current consumption given income z (so the graph of g is current consumption's Engel curve). When the utility function \tilde{u} is differentiable, u will be differentiable in x and we may plot an *iso-marginal utility diagram*, *i.e.*, a diagram that depicts the iso-marginal utility curves $IMU(c) \equiv \{(z, x) \in Z \times X : D_x u(x, z) = c\}$ for $c \in \mathbb{R}$. See Figure 1.

The IMU-diagram may be interpreted as follows. Since $u(x, z) = \tilde{u}(x, (1+r)(z-x))$, the marginal utility $D_x u(x, z)$ measures the gain from substituting one unit of consumption tomorrow for $1+r$ units of consumption today. Since there is a trade-off involved, this gain may be positive, negative, or zero as illustrated by the IMU curves in Figure 1 (when marginal utility equals zero, current consumption is optimal given income and so the curve $IMU(0)$ is current consumption's Engel curve). Different points on an IMU curve tell us that the gain remains the same whether she is poor and consumes little today, or rich and consumes more. *How much* more consumption is required for the gain to remain the same, is precisely what quasi-concave and quasi-convex differences place conditions on: quasi-concave differences is equivalent to $D_x u(x, z)$ being quasi-concave (Lemma 1). That $D_x u(x, z)$ is quasi-concave in turn means that the "better marginal utility (MU) sets" $\{(x, z) \in X \times Z : D_x u(x, z) \geq c\}$ are convex. See Figure 1 where the convex better MU sets account for the flattening IMU curves. That IMU curves flatten out tells us that as the consumer becomes wealthier, ever more modest increases in current consumption are needed to maintain a constant gain from substituting $1+r$ units today for 1 unit tomorrow. In the quasi-convex differences case, the IMU curves instead become steeper and steeper, hence ever larger consumption boosts are required for the gain not to change.

¹⁸For more general conditions on $D_x U(x, \bar{z}, z)$ that imply quasi-concavity of $D_x u(x, z)$, see Quah and Strulovici (2012).

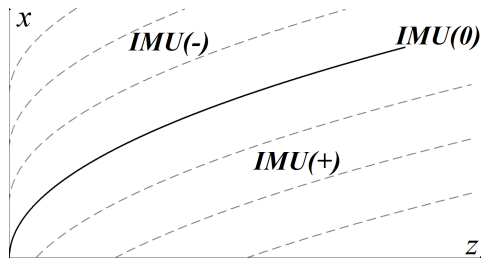


Figure 1: The zero IMU curve is current consumption's Engel curve.

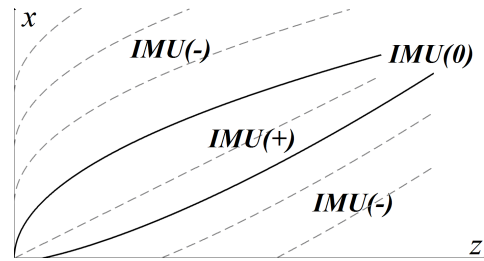


Figure 2: The upper zero IMU curve is current consumption's Engel curve.

Iso-Marginal Utility diagrams

Imagine now that the consumer receives an additional unit of income. Moreover, imagine that she spends a proportion of this on current consumption which precisely equals the proportion she spent on current consumption before. Then she will not be increasing current consumption “ever more modestly” and the marginal utility must consequently fall. To avoid this she must spend less on consumption today — hence when IMU curves flatten out, she must be spending a *smaller* proportion of the additional income on current consumption than the proportion she spent on it before.¹⁹ Behaviorally, flattening IMU curves means that current consumption becomes less and less effective in satisfying her needs as income increases: when income is low it is critical for her but she “tires” of it as consumption increases with income. This is of course a familiar phenomenon to all of us, and it applies not only as a realistic description of current versus future consumption (for most of us), but to a wide variety of physical goods. A good whose demand elasticity lies between 0 and 1 is called a *necessity good* (Varian (1992), p.117). The textbook example is potatoes; but if u exhibits quasi-concave differences then consumption today is a necessity good too.²⁰ So the previous interpretation of quasi-concave differences may be viewed as an interpreta-

¹⁹The same point can be made by instead considering the income margin, *i.e.*, the increase in current consumption associated with the last dollar received before the additional income. Graphically, this means that we pick a point on an IMU curve and move to the right along the IMU curve's tangent, in particular the proportion spent at the margin equals the slope of the IMU curve. Since the curve flattens (convexity), the tangent will lie above the initial point and so if the consumer follows the proportion at the margin before the income increase, she moves “above” the IMU curve and her marginal utility must fall. Since the proportion at the margin decreases with income under quasi-concave differences, it follows in particular that she spends a larger proportion of any additional income on consumption tomorrow than the proportion she spent on it before receiving the additional income.

²⁰When $c(0) = 0$, a concave consumption function implies that the good is a necessity good as seen by taking $y = 0$ in concavity's (differentiable) definition $c(y) \leq c(z) + c'(z)(y - z)$. But as this definition of concavity also shows, there is in general no firm relationship between elasticities and concavity of a policy function. It is thus deliberate that quasi-concave differences is not interpreted via elasticities in this section.

tion at the preference/utility level of what a necessity good is (in contrast, the standard definition is at the observable level of demand functions).

The conclusion that under quasi-concave differences, the decision variable becomes a less and less effective means of increasing the agent's payoff as the exogenous parameter z increases, is the key to understanding quasi-concave differences in general. Once we understand this it is easy to see why the MPC of Section 2.1 has a concave policy function when the marginal loss function is convex. To make this crystal clear imagine that the Lucas supply curve is given by $y = a_1 - a_2x + a_2z$ and $\pi = b_1 - b_2x$ where all constants are positive.²¹ In words, an increase in z is expansive without being inflationary and an interest rate increase moves the economy towards the origin of the Lucas supply curve. In this situation, counter-cyclical monetary policy (increasing x when z increases) is much like a "necessity good": At first, *i.e.*, when z increases from a low level, increasing the interest rate will both shift y and π towards targets (assuming that we begin above those targets) and so is a highly effective means of increasing the MPC's payoff. But as z keeps growing, decreasing efficiency kicks in because increasing x to force output back towards target simultaneously forces inflation further and further *below* target. This results in an increasing marginal payoff loss because the marginal loss function is convex.

Now, in each of the previous cases the policy function is increasing (in addition to being concave). But convexity of the better MU sets also captures concavity of g when g is not increasing. If we instead consider a decreasing policy function (*e.g.*, an inferior good), the interpretation is the same although the language changes slightly. IMU curves now become *steeper* corresponding to a decision variable that becomes increasingly ineffective at increasing the agent's payoff as the exogenous parameter z increases. As is clear, convexity of the better MU sets remains the critical feature behind a concave policy function. Similarly for non-monotonic functions where the previous terminologies can only be applied locally, but convexity of the better MU sets once again drives concavity.²²

All of the above is straight-forward once we see it in an IMU diagram. What is perhaps not as obvious is that quasi-concavity (or concavity) of u has nothing to do with the story. In Figure 1 the zero IMU's better set (denoted $IMU(+)$) is the entire set below the zero IMU. To be sure, this means that the payoff/utility function is quasi-concave in x since for fixed z it tells us that $u(x, z)$ is first increasing and then decreasing in x . But consider now Figure 2 where for fixed z , $u(x, z)$ is first decreasing, then increasing, and then again decreasing in x ; and so u is *not* quasi-concave in x . Since the better MU sets are convex, u exhibits quasi-concave differences. The zero IMU "curve" is now a correspondence consisting of two zero IMU curves. Since $u(x, z)$ is decreasing in x below the lower curve and increasing above it, any point on the lower zero-IMU curve is a local minimum. The upper zero IMU curve thus depicts the maxima so, just as in Figure 1, concavity of

²¹For this reduced form specification to make sense, target/natural levels must obviously be fixed as must any expectation not captured by z . See also footnote 4 for more details on the Lucas supply curve.

²²With quasi-convex differences the interpretation is in each instances reversed. Think of giving part of income to charity. Concave better IMU sets (quasi-convex differences) then means that as your income increases, you need to give progressively more of your current income to experience the same "warm glow" (utility gain). Presumably this would be because you need to feel you are making a sufficient sacrifice to get the same marginal utility effect, and you therefore have to progressively give more as a proportion of income for the sacrifice to keep its bite.

the consumption function is seen to obtain. Again, the reason is that the better MU sets are convex (quasi-concave differences) although now $IMU(+)$ is the lens between the $IMU(0)$ curves and not everything below $IMU(0)$ as in the quasi-concave setting of Figure 1.

To sum up, quasi-concave differences implies concavity of policy functions whether or not the policy function is monotone and regardless of any concavity or quasi-concavity assumptions. In this light, this Section's main result (Theorem 1 below) will come as no surprise to the reader. In fact, the proof below is just a formalization of the previous graphical argument that avoids using differentiability. There is only one complication related to solutions $g(z)$ touching the lower boundary of X , *i.e.*, solutions such that $g(z') = \inf X$ for some $z' \in Z$. In fact, such solutions will ruin any hope of obtaining a concave policy function for reasons that are easily seen graphically.

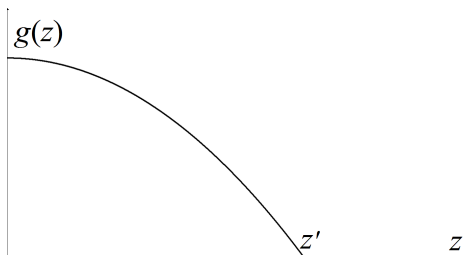


Figure 3: Concavity is destroyed when the policy function touches the lower boundary $\inf X = 0$.

In Figure 3, we see a policy function which at z' touches the lower boundary $\inf X = 0$ of the constraint set $X = \mathbb{R}_+$, and stays at this lower boundary point as z is further increased. It is evident that the resulting policy function will *not* be concave, even though it is concave for $z \leq z'$. As discussed at length in the working paper version of this paper (Jensen (2012)), this observation is robust: concave policy functions and lower boundary optimizers cannot coexist save for some *very* pathological cases. Of course, there is no problem if the optimization problem is unconstrained below, *i.e.*, if $\inf X = -\infty$. Nor is there a problem if we focus on the part of Z where the optimizer is interior (witness Figure 3 where we do have concavity when z is below z').

Theorem 1 (Concavity of the Policy Function) Let Z be a convex subset of a real vector space, and $X \subseteq \mathbb{R}$ a convex subset of the reals. Assume that the decision problem $\max_{x \in X} u(x, z)$ has a unique solution $g(z) = \arg \max_{x \in X} u(x, z) > \inf X$ for all $z \in Z$. Then if $u : X \times Z \rightarrow \mathbb{R}$ exhibits quasi-concave differences, $g : Z \rightarrow X$ is concave.

Proof. To simplify notation we set $\inf X = 0$ throughout. Pick arbitrary $z_1, z_2 \in Z$ and $\lambda \in [0, 1]$. Let $x_1 = g(z_1)$ and $x_2 = g(z_2)$ be the optimal decisions and define $x_\lambda = \lambda x_1 + (1 - \lambda)x_2$ and $z_\lambda = \lambda z_1 + (1 - \lambda)z_2$. Note that if $\{0\} \in \arg \max_{x \in [0, x_\lambda]} u(x, z_\lambda)$, then we necessarily have that $g(z_\lambda) \geq x_\lambda$ because 0 is not optimal by assumption. Since $g(z_\lambda) \geq x_\lambda \Leftrightarrow \lambda g(z_1) + (1 - \lambda)g(z_2) = x_\lambda \leq g(z_\lambda) = g(\lambda z_1 + (1 - \lambda)z_2)$, this is the same as saying that g is concave. We are now going to show that if $\{0\} \notin \arg \max_{x \in [0, x_\lambda]} u(x, z_\lambda)$, then

$$(9) \quad x_\lambda \in \arg \max_{x \in [0, x_\lambda]} u(x, z_\lambda).$$

Since (9) immediately implies that $\arg\max_{x \in X} u(x, z_\lambda) = g(z_\lambda) \geq x_\lambda$, which means that g is concave, Theorem 1 follows. Assume, by way of contradiction, that $x^* < x_\lambda$, where x^* is the largest point in $[0, x_\lambda]$ that maximizes $u(x, z_\lambda)$. Since 0 is not optimal, $x^* > 0$ and so we may choose a $\delta > 0$ sufficient small so that $x^* - \delta > 0$ and $x^* + \delta < x_\lambda$. Because $x_1, x_2 > 0$, and therefore $u(x_1, z_1) \geq u(x_1 - \delta, z_1)$ and $u(x_2, z_2) \geq u(x_2 - \delta, z_2)$, quasi-concave differences implies that:

$$(10) \quad u(x_\lambda, z_\lambda) - u(x_\lambda - \delta, z_\lambda) \geq \min\{u(x_1, z_1) - u(x_1 - \delta, z_1), u(x_2, z_2) - u(x_2 - \delta, z_2)\} \geq 0.$$

Since x^* is optimal,

$$(11) \quad u(x^*, z_\lambda) - u(x^* - \delta, z_\lambda) \geq 0.$$

And since x^* is the largest such optimal point in $[0, x_\lambda]$ and $x^* + \delta < x_\lambda$,

$$(12) \quad u(x^* + \delta, z_\lambda) - u(x^*, z_\lambda) < 0.$$

Because u exhibits quasi-concave differences, $u(x, z_\lambda) - u(x - \delta, z_\lambda)$ is quasi-concave in x for any δ sufficiently close to zero. It follows therefore from (11)-(12) that,

$$(13) \quad u(x_\lambda, z_\lambda) - u(x_\lambda - \delta, z_\lambda) < 0.$$

But this contradicts (10).

■

Corollary 1 (Convexity of the Policy Function) Let Z be a convex subset of a real vector space, and $X \subseteq \mathbb{R}$ a convex subset of the reals. Assume that the decision problem $\max_{x \in X} u(x, z)$ has a unique solution $g(z) = \arg\max_{x \in X} u(x, z) < \sup X$ for all $z \in Z$. Then if $u : X \times Z \rightarrow \mathbb{R}$ exhibits quasi-convex differences, $g : Z \rightarrow X$ is convex.

Proof. Let $-X \equiv \{-x \in \mathbb{R} : x \in X\}$. Apply Theorem 1 to the optimization problem $\max_{\tilde{x} \in -X} u(-\tilde{x}, z)$ and use that the policy function of this problem is concave if and only if g is convex. ■

We end this section with a discussion of the relationship between (strict) quasi-concavity of $u(\cdot, z)$ and quasi-concave differences.²³ To simplify, focus is on the smooth case (but the conclusions are true in general as the reader may easily verify). $u(x, z)$ is strictly quasi-concave in x if and only if the partial derivative $D_x u(x, z)$ is strictly positive on $(\inf X, x^*)$ and strictly negative on $(x^*, \sup X)$ (here x^* may equal the infimum or supremum of X , *i.e.*, the function may be monotone). By Lemma 1 follows that if u exhibits quasi-concave differences then $D_x u(x, z)$ is quasi-concave in x . Hence $D_x u(x, z)$ is increasing on $(\inf X, \tilde{x})$ and decreasing on $(\tilde{x}, \sup X)$ for some \tilde{x} (again \tilde{x} may be equal to the supremum or infimum in which case the first partial derivative will be monotone). Comparing the respective conditions on $D_x u(x, z)$ it is clear that there can be no direct relationship between (strict) quasi-concavity of the objective function and quasi-concave differences: Quasi-concave differences is fully compatible with $u(\cdot, z)$ being first decreasing, then increasing, and then decreasing again (see Figure 2). Hence quasi-concave differences does not

²³This discussion is included on the request by a referee.

imply quasi-concavity. On the other hand, a function whose first derivative is always positive but first strictly decreases and then strictly increases will be strictly quasi-concave but cannot exhibit quasi-concave differences. The one similarity I am aware of between the conditions relates to the structure of the set of optimizers. A quasi-concave function always has a convex set of maximizers. If $D_x u(x, z)$ is quasi-concave in x and the infimum $\inf X$ is not a maximizer as assumed in Theorem 1, one similarly sees that the set of maximizers must be convex (this is very easy to see in the smooth case, but it is true in general). If we define *strictly quasi-concave differences* as in Definition 1 by replacing the word quasi-concave with strictly quasi-concave, it will follow by the same line of reasoning that there can be at most one maximizer. This strict version of quasi-concave differences thus parallel's strict quasi-concavity in securing a unique optimizer — but as was just explained the conditions are logically independent of each other (neither implies the other).

3.2 The General Case

In applications such as the income allocation model of Section 2.2, it is too restrictive to assume that the constraint set X is fixed. Further, one may face decision problems with multiple solutions unless strict quasi-concavity in x or some similar condition holds. We then face the general decision problem

$$(14) \quad G(z) = \arg \max_{x \in \Gamma(z)} u(x, z).$$

Here $\Gamma : Z \rightarrow 2^X$ is the *constraint correspondence* and $G : Z \rightarrow 2^X$ is the *policy correspondence*. A policy function is now a selection from G , *i.e.*, a function $g : Z \rightarrow X$ with $g(z) \in G(z)$ for all $z \in Z$. The assumption of a one-dimensional decision variable $X \subseteq \mathbb{R}$ is maintained to keep things simple. Appendix C deals with the general case where X is a subset of a topological vector lattice.

For (14), the result of Topkis (1978) tells us that if u exhibits increasing differences and Γ is an ascending correspondence, then G is ascending. The precise definition of an ascending correspondence is not important for us here; it suffices to say that it naturally extends the notion that a function is increasing to a correspondence. As it turns out, the conclusion of Theorem 1 generalizes in a very similar manner. The only question is how to extend concavity/convexity from a function to a correspondence in a suitable way for our results.²⁴

Definition 3 (Concave Correspondences) A correspondence $\Gamma : Z \rightarrow 2^X$ is *concave* if for all $z_1, z_2 \in Z$, $x_1 \in \Gamma(z_1)$, $x_2 \in \Gamma(z_2)$, and $\lambda \in [0, 1]$, there exists $x \in \Gamma(\lambda z_1 + (1 - \lambda)z_2)$ with $x \geq \lambda x_1 + (1 - \lambda)x_2$.

For illustrations, see Figures 1-2 where the sets $IMU(+)$ depict graphs of concave correspondences. In parallel with concave/convex functions, $\Gamma : Z \rightarrow 2^X$ is said to be *convex* if $-\Gamma : Z \rightarrow 2^{-X}$ is concave where $-\Gamma(z) \equiv \{-x \in \mathbb{R} : x \in \Gamma(z)\}$. Definition 3 naturally generalizes concavity of a function to a correspondence. In particular, one immediately sees that if Γ is single-valued, then it is concave if and only if the function it defines is concave (and similarly, Γ is convex if and only if the function it defines is convex). Recall that a correspondence $\Gamma : Z \rightarrow 2^X$ has a *convex graph*

²⁴The following definition can be found in Kuroiwa (1996) who also offers an extensive discussion of set-valued convexity. Lemma 2 below appears to be new, however.

if $\{(x, z) \in X \times Z : x \in \Gamma(z)\}$ is a convex subset of $X \times Z$. Convexity of a correspondence's graph is a much stronger requirement than concavity and convexity of Γ . In fact, a correspondence with a convex graph is both concave and convex.²⁵ Furthermore, if Γ has a convex graph, it also has convex values, *i.e.*, $\Gamma(z)$ must be a convex subset of X for all $z \in Z$. In contrast, convex values is not implied by either concavity or convexity and so will have to be assumed directly when needed (as it will be below).

The following result sheds further light on the definition. It tells us that concavity and convexity of a correspondence is intimately tied to concavity and convexity of extremum selections (when they exist). For most applications, this result is also enough to establish that a given constraint correspondence is concave or convex since it covers inequality constraints where $\Gamma(z) = \{x \in X : \underline{\gamma}(z) \leq x \leq \bar{\gamma}(z)\}$.

Lemma 2 (Extremum Selection Criteria) If $\Gamma : Z \rightarrow 2^X$ admits a greatest selection, $\bar{\gamma}(z) \equiv \sup \Gamma(z) \in \Gamma(z)$ for all $z \in Z$, then Γ is concave if and only if $\bar{\gamma} : Z \rightarrow X$ is a concave function. Likewise, if Γ admits a least selection $\underline{\gamma}(z) \equiv \inf \Gamma(z) \in \Gamma(z)$ all $z \in Z$, Γ is convex if and only if $\underline{\gamma}$ is a convex function.

Proof. Only the concave case is proved. Since Γ is concave, we will for any $z_1, z_2 \in Z$, and $\lambda \in [0, 1]$ have an $x \in \Gamma(\lambda z_1 + (1-\lambda)z_2)$ with $x \geq \lambda \underline{\gamma}(z_1) + (1-\lambda)\underline{\gamma}(z_2)$. Since $\bar{\gamma}(\lambda z_1 + (1-\lambda)z_2) \geq x$, $\bar{\gamma}$ is concave. To prove the converse, pick $z_1, z_2 \in Z$ and $x_1 \in \Gamma(z_1)$, $x_2 \in \Gamma(z_2)$. Since the greatest selection is concave, $x = \bar{\gamma}(\lambda z_1 + (1-\lambda)z_2) \geq \lambda \bar{\gamma}(z_1) + (1-\lambda)\bar{\gamma}(z_2) \geq \lambda x_1 + (1-\lambda)x_2$. Since $x \in \Gamma(\lambda z_1 + (1-\lambda)z_2)$, Γ is concave. ■

Often one is able to spot a concave correspondence immediately from this Lemma. For example, we see that the union of the zero IMU curves in Figure 2 is the graph of a concave correspondence where the lower zero IMU is the least selection and the upper zero IMO the greatest selection. Lemma 2 also shows exactly how concavity and convexity relates to other known convexity concepts for correspondences. In a diagram with z on the first axis and x on the second axis, draw the graph of a concave function $\bar{\gamma}$. Now extend this graph to the graph of a correspondence Γ by drawing freely anything at or below the graph of $\bar{\gamma}$. Then the resulting correspondence is concave by Lemma 2. So in Figure 2 we could draw anything below the upper zero IMU and would still have a concave correspondence. As an aside, it is evident from these Figures that convex values as well as a convex graph are not implied — and it is equally evident that a convex graph implies that the correspondence is both concave and convex (these facts were discussed in a more technical manner a moment ago).

Theorem 2 (Concavity of the Policy Correspondence) Let Z be a convex subset of a real vector space, and $X \subseteq \mathbb{R}$ a convex subset of the reals. Assume that the decision problem $\max_{x \in \Gamma(z)} u(x, z)$ has solution $G(z) = \arg \max_{x \in \Gamma(z)} u(x, z) \neq \emptyset$ for all $z \in Z$ and that the infimum of $\Gamma(z)$ is never optimal, $x \in G(z) \Rightarrow x > \inf \Gamma(z)$. Then if $u : X \times Z \rightarrow \mathbb{R}$ exhibits quasi-concave differences and Γ is concave and has convex values, $G : Z \rightarrow 2^X$ is concave.

²⁵To see this, simply pick $x = \lambda x_1 + (1-\lambda)x_2 \in \Gamma(\lambda z_1 + (1-\lambda)z_2)$ in Definition 3.

Proof. Pick any $z_1, z_2 \in Z$, $x_1 \in G(z_1)$, and $x_2 \in G(z_2)$. Setting $x_\lambda = \lambda x_1 + (1 - \lambda)x_2$, and $z_\lambda = \lambda z_1 + (1 - \lambda)z_2$, we must show that there exists $\hat{x} \in G(z_\lambda)$ with $\hat{x} \geq x_\lambda$. As in the proof of Theorem 1, we use quasi-concave differences to conclude that $u(x_\lambda, z_\lambda) - u(x_\lambda - \delta, z_\lambda) \geq 0$ for any sufficiently small $\delta > 0$. We are clearly done if there does not exist $x \in \Gamma(z_\lambda)$ with $x < x_\lambda$. So assume that such an x exists. By concavity of Γ , there also exists $\tilde{x} \in \Gamma(z_\lambda)$ with $\tilde{x} \geq x_\lambda$. Since Γ has convex values therefore $[x, \tilde{x}] \subseteq \Gamma(z_\lambda)$. Since $x_\lambda \in [x, \tilde{x}]$ and $\inf \Gamma(z_\lambda)$ cannot be optimal, we may now proceed precisely as in the proof of Theorem 1 and use quasi-concavity of $u(x, z_\lambda) - u(x - \delta, z_\lambda)$ in x to show that there must exist a $\hat{x} \in \Gamma(z_\lambda)$ with $\hat{x} \geq x_\lambda$. ■

Corollary 2 (Convexity of the Policy Correspondence) Let Z be a convex subset of a real vector space, and $X \subseteq \mathbb{R}$ a convex subset of the reals. Assume that the decision problem $\max_{x \in \Gamma(z)} u(x, z)$ has solution $G(z) = \arg \max_{x \in \Gamma(z)} u(x, z) \neq \emptyset$ for all $z \in Z$ and that the supremum of $\Gamma(z)$ is never optimal, $x \in G(z) \Rightarrow x < \sup \Gamma(z)$. Then if $u : X \times Z \rightarrow \mathbb{R}$ exhibits quasi-convex differences and Γ is convex and has convex values, $G : Z \rightarrow 2^X$ is convex.

Proof. Let $-\Gamma(z) \equiv \{-x \in \mathbb{R} : x \in \Gamma(z)\}$. Apply Theorem 2 to the optimization problem $\max_{\tilde{x} \in -\Gamma(z)} u(-\tilde{x}, z)$ and use that the policy correspondence of this problem is concave if and only if G is convex. ■

If the conditions of Theorem 2 hold and the policy correspondence is single-valued $G = \{g\}$, then g must be a concave function by Lemma 2. Hence Theorem 1 is a special case of Theorem 2. From Lemma 2 also follows that when u is upper semi-continuous and Γ has compact values so that G has a greatest selection, this greatest selection must be concave.²⁶ Finally, note that just like in the theory of monotone comparative statics (Topkis (1978), Milgrom and Shannon (1994)), these observations are valid without assuming that the objective function is quasi-concave in the decision variable.

3.3 A User's Guide and the Model of Melitz (2003)

This subsection provides a practitioners' guide to Theorems 1 and 2. We begin by noting an immediate consequence of Lemma 1:

Lemma 3 (Quasi-Concave Differences for Thrice Differentiable Functions) *A thrice differentiable function $u : X \times Z \rightarrow \mathbb{R}$ where $X, Z \subseteq \mathbb{R}$ exhibits quasi-concave differences if and only if*

$$(15) 2D_{xx}^2 u(x, z) D_{xz}^2 u(x, z) D_{xxz}^3 u(x, z) \geq [D_{xx}^2 u(x, z)]^2 D_{xzz}^3 u(x, z) + [D_{xz}^2 u(x, z)]^2 D_{xxx}^3 u(x, z).$$

Proof. (15) is the non-negative bordered Hessian criterion for quasi-concavity of $D_x u(x, z)$ (see e.g. Mas-Colell et al (1995), pp.938-939). By Lemma 1, this is equivalent to $u(x, z)$ exhibiting quasi-concave differences. ■

²⁶Note that it is unreasonable to expect the *least* selection to be concave also. In fact, this would not characterize any reasonable concavity-type condition for a correspondence (in the case of a correspondence with a convex graph, for example, the greatest selection is concave and the least selection is convex).

What is nice about condition (15) is that when the payoff function is smooth, it makes the verification of quasi-concave differences completely tractable.

We next show in a step-by-step manner how one goes about using the results of the previous pages. In particular, we will be using (15). While the step-by-step structure of the argument is entirely general, we focus for concreteness on a well known model from the international trade literature due to Melitz (2003).

Each firm in a continuum $[0, 1]$ chooses output $x \geq 0$ in order to maximize profits. A firm with cost parameter $z > 0$, can produce x units of the output by employing $l = zx + f$ workers where $f > 0$ is a fixed overhead (Melitz (2003), p.1699).²⁷ The frequency distribution of the cost parameter z across the firms is η_z . With revenue function R , a firm with cost parameter z chooses $x \geq 0$ in order to maximize

$$(16) \quad u(x, z) = R(x) - zx - f.$$

Let $G(z) = \operatorname{argmax}_{x \geq 0} [R(x) - zx - f]$ denote the optimal output(s) given z . To show that G is concave or convex, we may apply Theorem 1 or the more general Theorem 2. For Theorem 1, $R(x) - zx - f$ must be strictly quasi-concave or satisfy some other condition that guarantees that firms have a unique optimal output level, $G(z) = \{g(z)\}$. As one easily verifies, (15) holds if and only if

$$(17) \quad R''' \leq 0.$$

So by Theorem 1, if the revenue function has a non-positive third derivative and optimizers are unique, the policy function g is concave on $z \in \{z \in Z : g(z) > 0\}$, *i.e.*, when attention is restricted to the set of active firms. It follows from Observation 2 on page 5 that aggregate output $\int_{[0,1]} g(z) \eta_z(dz)$ decreases when firms become more diverse (a mean preserving spread to the distribution η_z). Such “decreasing returns to diversity” is easily understood in light of (17) which says that the marginal revenue function is concave. A concave marginal revenue function tells us that if we consider two firms that produce x_1 and x_2 , respectively, and the firms have different productivities $z_1 \neq z_2$, then the average of their marginal revenues $(R'(x_1) + R'(x_2))/2$ will be *lower* than the marginal revenue of a firm which produces the average output $(x_1 + x_2)/2$ and has the average productivity $(z_1 + z_2)/2$. In particular, when x_1 and x_2 are optimal for the respective firms and the marginal revenues therefore equal zero, the marginal revenue of the average productivity firm will be positive, and it is therefore optimal for it to produce more than the average $(x_1 + x_2)/2$. Putting these observations together, we see that if production is spread across diverse firms, then total output is lower than if production takes place among an equal number of more similar firms. For an extensive discussion of concave marginal revenue functions in the theory of production, the reader is referred to Leahy and Whited (1996).

By appealing to Observations 1-5 on page 5 one can further go on to predict how the distribution of the firms’ outputs changes when η_z is subjected to mean-preserving spreads or other stochastic order changes. For example, the distribution of the outputs of a more diverse set of firms will be second-order stochastically dominated by the distribution of outputs of a less diverse set

²⁷In terms of Melitz’ notation, z is the inverse of the firm’s productivity level φ .

of firms when (17) holds (this is again by Observation 2). If u exhibits quasi-convex differences (reverse the inequality (15) yielding the condition $R''' \geq 0$), we instead get a convex policy function by Corollary 1; there is “increasing returns to diversity”, and all of the previous conclusions are reversed.

The limitations of Theorem 1 are evident in the current situation since monopolistically competitive firms’ objectives will often not be strictly quasi-concave or even quasi-concave and so uniqueness of the optimizer is difficult to ensure in general. One can then instead use Theorem 2. Assuming only that solutions exist (so that G is well-defined), we can conclude that G is a concave correspondence when the revenue function R has a non-positive third derivative. Hence the greatest selection from the policy correspondence will be a concave function (Lemma 2), and the maximum aggregate output decreases with a mean-preserving spread to η_z .²⁸

4 Bayesian Games

The purpose of this section is to use Theorem 1 to deal with increased uncertainty in Bayesian games. In Section 2.1, the central bank (the MPC) was the only agent who made a decision, and that decision was not influenced by the other agent’s action (the forecast). The increase in uncertainty on the other hand, affected only the forecaster (by assumption the MPC knew its private signal). In reality, central banks take other agents’ responses into account when setting interest rates — and those agents will be aware of this and take *that* into account. So, any increase in uncertainty spills over between the agents. Such considerations naturally lead us to study distributional comparative statics in Bayesian games.

A Bayesian game consists of a set of players $\mathcal{I} = \{1, \dots, I\}$, taken here to be finite, where player $i \in \mathcal{I}$ receives a private signal $z_i \in Z_i \subseteq \mathbb{R}$ drawn from a distribution μ_{z_i} on $(Z_i, \mathcal{B}(Z_i))$.²⁹ Agents maximize their objectives and a Bayesian equilibrium is just a Nash equilibrium of the resulting game (defined precisely in a moment). The question we ask is this: How will the set of equilibria be affected if one or more signal distributions μ_{z_i} are subjected to mean-preserving spreads or second-order stochastic dominance shifts? If $i = 1$ is the MPC and z_1 represents the MPC’s assessment of the Lucas supply curve, we are back in the setting of Section 2.1 except that we now allow for the game-theoretic interaction between the MPC and the other agents in the economy.

Assuming that private signals are independently distributed, an *optimal strategy* is a measurable mapping $g_i : Z_i \rightarrow X_i$ such that for almost every $z_i \in Z_i$,

$$(18) \quad g_i(z_i) \in \arg \max_{x_i \in X_i} \int_{z_{-i} \in Z_{-i}} U_i(x_i, g_{-i}(z_{-i}), z_i) \mu_{z_{-i}}(dz_{-i}).$$

Here $X_i \subseteq \mathbb{R}$ is agent i ’s action set and $g_{-i} = (g_j)_{j \neq i}$ are the strategies of the opponents.

²⁸If R is assumed to be strictly concave so that $R'' < 0$, we might alternatively have applied the implicit function theorem (IFT) to the first-order condition $R'(x) - z = 0$. This yields $x = g(z)$ where $D_z g = R''$. The IFT is particularly easy to use in this case, and we immediately see that (17) once again ensures concavity of g (this is because $D_z g$ decreasing precisely means that g is concave). Note, however, that Theorem 2 (and also Theorem 1) applies to many situations which the IFT is unable to address.

²⁹ $\mathcal{B}(\cdot)$ denotes the Borel subsets of a given set.

Assumption 1 For every i , the optimal strategy g_i exists and is unique.

This assumption is satisfied if X_i is compact and $U_i(x_i, x_{-i}, z_i)$ is strictly concave in x_i (risk aversion), and continuous in (x_i, x_{-i}, z_i) . If uniqueness is not assumed, we could instead use Theorem 2 but we favor here simplicity over generality. Note that what we here call an optimal strategy is very closely related to policy functions (specifically, $g_i : Z_i \rightarrow X_i$ is a policy function when it satisfies (18) for all $z_i \in Z_i$).

A *Bayesian equilibrium* is a strategy profile $g^* = (g_1^*, \dots, g_I^*)$ such that for each player i , $g_i^* : Z_i \rightarrow X_i$ is an optimal strategy given the opponents' strategies $g_{-i}^* : Z_{-i} \rightarrow X_{-i}$. The *optimal distribution* of an agent i is the measure on $(X_i, \mathcal{B}(X_i))$ given by:

$$(19) \quad \mu_{x_i}(A) = \mu_{z_i} \{z_i \in Z_i : g_i(z_i) \in A\}, \quad A \in \mathcal{B}(X_i)$$

For given opponents' strategies the decision problem in (18) coincides with the MPC's objective in Section 2.1. In particular, we know how changes in μ_{z_i} affect the optimal distribution of the player μ_{x_i} when g_i is concave or convex (see Observations 1-5 on page 5). Combining with Theorem 1 we immediately get:

Lemma 4 Consider a player $i \in \mathcal{I}$ and let Assumption 1 be satisfied.

1. If $\int_{z_{-i} \in Z_{-i}} U_i(x_i, g_{-i}(z_{-i}), z_i) \mu_{z_{-i}}(dz_{-i})$ exhibits quasi-concave differences in x_i and z_i , and no element on the lower boundary of X_i ($\inf X_i$) is optimal, then a mean-preserving spread to μ_{z_i} will lead to a second-order stochastic dominance decrease in the optimal distribution μ_{x_i} .
2. If $\int_{z_{-i} \in Z_{-i}} U_i(x_i, g_{-i}(z_{-i}), z_i) \mu_{z_{-i}}(dz_{-i})$ exhibits quasi-convex differences in x_i and z_i , and no element on the upper boundary of X_i ($\sup X_i$) is optimal, then a mean-preserving spread to μ_{z_i} will lead to a convex-increasing order increase in the optimal distribution μ_{x_i} .

As we saw in Section 3.1, the conditions of Lemma 4 are easy to verify if U_i is differentiable in x_i (see also the main Theorem below which provides sufficient conditions in the differentiable case). Note that whether increased uncertainty increases or reduces mean actions depends on whether the payoff function exhibits quasi-convex or quasi-concave differences. But in either case, the optimal distributions will be more variable which will spill over to the other players and make *everybody's* game environments more uncertain. To deal with this, we need the following straightforward generalization of a result found in Rothschild and Stiglitz (1971).³⁰

Lemma 5 Let Assumption 1 be satisfied and let $g_i(z_i, \mu_{x_{-i}}) = \arg \max_{x_i \in X_i} \int U_i(x_i, x_{-i}, z_i) \mu_{x_{-i}}(dx_{-i})$. Then for $j \neq i$:

³⁰Rothschild and Stiglitz (1971) consider mean-preserving spreads in the differentiable case. If U_i is differentiable in x_i , the main condition of Lemma 5 is equivalent to the concavity of $D_{x_i} U_i(x_i, x_{-i}, \cdot)$ which exactly is the assumption of Rothschild and Stiglitz (1971). See also Athey (2002) for related results. It is by combining such (known) results with this paper's new results that we are able to make progress on DCS in Bayesian games.

1. If $U_i(\tilde{x}_i, x_{-i}, z_i) - U_i(x_i, x_{-i}, z_i)$ is concave in x_j for all $\tilde{x}_i \geq x_i$, then $g_i(z_i, \mu_{x_{-i-j}}, \tilde{\mu}_{x_j}) \leq g_i(z_i, \mu_{x_{-i}})$ whenever $\tilde{\mu}_{x_j}$ is a mean-preserving spread of μ_{x_j} .
2. If $U_i(\tilde{x}_i, x_{-i}, z_i) - U_i(x_i, x_{-i}, z_i)$ is concave and increasing in x_j for all $\tilde{x}_i \geq x_i$, then $g_i(z_i, \mu_{x_{-i}}) \leq g_i(z_i, \mu_{x_{-i-j}}, \tilde{\mu}_{x_j})$ whenever $\tilde{\mu}_{x_j}$ second-order stochastically dominates μ_{x_j} .

If in these statements concavity in x_j is replaced with convexity, the first conclusion changes to: $g_i(z_i, \mu_{x_{-i-j}}, \tilde{\mu}_{x_j}) \geq g_i(z_i, \mu_{x_{-i}})$ whenever $\tilde{\mu}_j$ is a mean-preserving spread of μ_{x_j} ; and the second conclusion changes to $g_i(z_i, \mu_{x_{-i}}) \leq g_i(z_i, \mu_{x_{-i-j}}, \tilde{\mu}_{x_j})$ whenever $\tilde{\mu}_{x_j}$ dominates μ_{x_j} in the convex-increasing order.

Proof. Statement 1 is a direct application of Topkis' theorem (Topkis (1978)) which in the situation with a one-dimensional decision variable and unique optimizers says that the optimal decision will be non-decreasing [non-increasing] in parameters if the objective exhibits increasing differences [decreasing differences]. The conclusion thus follows from the fact that $\int U_i(x_i, x_{-i}, z_i) \mu_{x_{-i}}(dx_{-i})$ exhibits decreasing differences in x_i (with the usual order) and μ_{x_j} (with the mean-preserving spread order \succeq_{cx}) if and only if the assumption of the statement holds. Also by Topkis' theorem, if $U_i(\tilde{x}_i, x_{-i}, z_i) - U_i(x_i, x_{-i}, z_i)$ is increasing in x_j for $j \neq i$ and for all $\tilde{x}_i \geq x_i$, then $\tilde{\mu}_{x_j} \succeq_{st} \mu_{x_j} \Rightarrow g_i(z_i, \mu_{x_{-i-j}}, \tilde{\mu}_{x_j}) \geq g_i(z_i, \mu_{x_{-i}})$ (here \succeq_{st} denotes the first-order stochastic dominance order). From this and Statement 1 follows Statement 2 because it is always possible to split a second order stochastic dominance increase \succeq_{cvi} into a mean preserving contraction \succeq_{cv} followed by a first order stochastic dominance increase (Formally, if $\tilde{\mu}_{x_j} \succeq_{cvi} \mu_{x_j}$, then there exists a distribution $\hat{\mu}_{x_j}$ such that $\tilde{\mu}_{x_j} \succeq_{st} \hat{\mu}_{x_j} \succeq_{cv} \mu_{x_j}$). The convex case is proved by a similar argument and is omitted. ■

For given distributions of private signals $\mu_z = (\mu_{z_1}, \dots, \mu_{z_l})$ let $\Phi(\mu_z)$ denote the *set of equilibrium distributions, i.e.*, the set of optimal distributions $\mu_x = (\mu_{x_1}, \dots, \mu_{x_l})$ where μ_{x_i} are given in (19) with $g_i = g_i^*$ and (g_1^*, \dots, g_l^*) is one of the (possibly many) Bayesian equilibria. Fix a given stochastic order \succeq on the probability space of optimal distributions and consider a shift in the distribution of private signals from μ_z to $\tilde{\mu}_z$. In the Theorem below, \succeq will be either the second-order stochastic dominance order or the convex-increasing order; and the shift from μ_z til $\tilde{\mu}_z$ will be a mean-preserving spread. The set of equilibrium distributions then *increases* in the order \succeq if

$$(20) \quad \forall \mu_x \in \Phi(\mu_z) \exists \tilde{\mu}_x \in \Phi(\tilde{\mu}_z) \text{ with } \tilde{\mu}_x \succeq \mu_x \text{ and } \forall \tilde{\mu}_x \in \Phi(\tilde{\mu}_z) \exists \mu_x \in \Phi(\mu_z) \text{ with } \tilde{\mu}_x \succeq \mu_x$$

If the order \succeq is reversed in (20), the set of equilibrium distributions *decreases*. If $\Phi(\mu_z)$ and $\Phi(\tilde{\mu}_z)$ have least and greatest elements, then (20) implies that the least element of $\Phi(\mu_z)$ will be smaller than the least element of $\Phi(\tilde{\mu}_z)$ and the greatest element of $\Phi(\mu_z)$ will be smaller than the least element of $\Phi(\tilde{\mu}_z)$ (Smithson (1971), Theorem 1.7). In particular, if the equilibria are unique and we therefore have a function ϕ such that $\Phi(\mu_z) = \{\phi(\mu_x)\}$ and $\Phi(\tilde{\mu}_z) = \{\phi(\tilde{\mu}_x)\}$, we get $\phi(\mu_z) \leq \phi(\tilde{\mu}_z)$ which simply means that the function ϕ is increasing.

Theorem 3 (Mean Preserving Spreads in Bayesian Games) Consider a Bayesian game as described above and let $\mu_z = (\mu_{z_i})_{i \in \mathcal{I}}$ and $\tilde{\mu}_z = (\tilde{\mu}_{z_i})_{i \in \mathcal{I}}$ be two distributions of private signals.³¹

1. Suppose all assumptions of Lemma 4.1 are satisfied and $U_i(\tilde{x}_i, x_{-i}, z_i) - U_i(x_i, x_{-i}, z_i)$ is increasing and concave in x_{-i} for all $\tilde{x}_i \geq x_i$. If $\tilde{\mu}_{z_i}$ is a mean-preserving spread of μ_{z_i} for any subset of the players, then the set of equilibrium distributions decreases in the second-order stochastic dominance order (in particular the agents' mean actions decrease).
2. Suppose all assumptions of Lemma 4.2 are satisfied and $U_i(\tilde{x}_i, x_{-i}, z_i) - U_i(x_i, x_{-i}, z_i)$ is increasing and convex in x_{-i} for all $\tilde{x}_i \geq x_i$. If $\tilde{\mu}_{z_i}$ is a mean-preserving spread of μ_{z_i} for any subset of the players, then the set of equilibrium distributions increases in the convex-increasing order (in particular the agents' mean actions increase).

If U_i is differentiable in x_i , all of these assumptions are satisfied if $D_{x_i} U_i(x_i, x_{-i}, z_i)$ is increasing in x_{-i} and either concave in (x_i, z_i) and x_{-i} [case 1] or convex in (x_i, z_i) and x_{-i} [case 2].

Proof. As in previous proofs, let \succeq_{cvi} denote the concave-increasing (second-order stochastic dominance) order and \succeq_{st} denote the first-order stochastic dominance order. Recast the game in terms of optimal distributions: Agent i 's problem is to find a measurable function g_i which for *a.e.* $z_i \in Z_i$ maximizes $\int_{x_{-i} \in X_{-i}} U_i(x_i, x_{-i}, z_i) \mu_{x_{-i}}(dx_{-i})$. The policy function $x_i = g_i(z_i, \mu_{x_{-i}})$ determines $\mu_{x_i}(A) = \mu_{z_i}\{z_i \in Z_i : g_i(z_i, \mu_{x_{-i}}) \in A\}$ ($A \in \mathcal{B}(X_i)$). An equilibrium is a vector $\mu_x^* = (\mu_{x_1}^*, \dots, \mu_{x_I}^*)$ such that for all $i \in \mathcal{I}$: $\mu_{x_i}^*(A) = \mu_{z_i}\{z_i \in Z_i : g_i(z_i, \mu_{x_{-i}}^*) \in A\}$, all $A \in \mathcal{B}(X_i)$. Letting $f_i(\mu_{x_{-i}}, \mu_{z_i})$ denote agent i 's optimal distribution given $\mu_{x_{-i}}$ and μ_{z_i} , an equilibrium is a fixed point of $f = (f_1, \dots, f_I)$. By 2 of Lemma 5 and Observation 1 on page 5, $\tilde{\mu}_{x_j} \succeq_{cvi} \mu_{x_j} \Rightarrow f_i(\mu_{x_{-i-j}}, \tilde{\mu}_{x_j}, \mu_{z_i}) \succeq_{st} f_i(\mu_{x_{-i}}, \mu_{z_i})$ for all $j \neq i$. Since first-order stochastic dominance implies second-order stochastic dominance, $f_i(\mu_{x_{-i-j}}, \tilde{\mu}_{x_j}, \mu_{z_i}) \succeq_{st} f_i(\mu_{x_{-i}}, \mu_{z_i}) \Rightarrow f_i(\mu_{x_{-i-j}}, \tilde{\mu}_{x_j}, \mu_{z_i}) \succeq_{cvi} f_i(\mu_{x_{-i}}, \mu_{z_i})$. It follows that the mapping f is monotone when μ_x 's underlying probability space is equipped with the product order \succeq_{cvi} . Again with the order \succeq_{cvi} on optimal distributions, it follows from Lemma 4 that each f_i is decreasing in μ_{z_i} with the convex (mean-preserving spread) order on μ_{z_i} 's underlying probability space. f will also be continuous (it is a composition of continuous functions) and so the theorem's conclusions follow directly from Theorem 3 in Acemoglu and Jensen (2015) (the conditions of that Theorem are immediately satisfied when f is viewed as a correspondence). For the second statement of the theorem the argument is precisely the same except that one now equips the set of optimal distributions with the convex-increasing order and notes that f_i is monotone when the private distributions μ_{z_i} 's underlying probability spaces are equipped with the mean preserving spread order. The differentiability conditions presented at the end of the theorem follow from Lemma 1. ■

Note that under the assumptions of Theorem 3, the game is supermodular. Under the additional conditions of the following corollary, the game is monotone (*i.e.*, a Bayesian game of strategic complementarities, see Van Zandt and Vives (2007)). The proof follows along the same lines as the proof of Theorem 3 and is omitted.

³¹In the following statements, it is to be understood that any distribution μ_{z_i} that is *not* replaced with a mean-preserving spread $\tilde{\mu}_{z_i}$ is kept fixed.

Corollary 3 (Second-Order Stochastic Dominance Changes) If in addition to the assumptions of Theorem 3, it is assumed that $U_i(\tilde{x}_i, x_{-i}, z_i) - U_i(x_i, x_{-i}, z_i)$ is increasing in z_i , then if $\tilde{\mu}_{z_i}$ second-order stochastically dominates μ_{z_i} for any subset of the players, the set of equilibria decreases in the second-order stochastic dominance order in case 1. In case 2., the set of equilibria increases in the convex-increasing order when $\tilde{\mu}_{z_i}$ dominates μ_{z_i} in the convex-increasing order for any subset of the players.

There are many interesting applications of Theorem 3, ranging from auction theory to the Diamond search model. Here we will consider two applications in turn beginning with the Bayesian extension of the model from Section 2.1.

4.1 Forecasting Monetary Policy (Continued)

In Section 2.1, \bar{z} in the expected utility formulation (3) was interpreted as an additional exogenous variable (the oil price for example). That formulation was chosen deliberately because it also covers Bayesian games where \bar{z} is now instead the private signal z_2 of a second agent (“the public”):

$$(21) \quad u_1(x_1, z_1) = \int U_1(x_1, g_2(z_2), z_1) \mu_{z_2}(dz_2).$$

z_1 is the MPC’s private signal, and $-U_1$ its loss function. As discussed in Section 2.1, the loss function incorporates the MPC’s assessment of the Lucas supply curve which depends on the MPC’s private signal z_1 but now also depends on what the public does, $g_2(z_2)$. What the public does is, in turn, uncertain to the MPC since it depends on the public’s private signal z_2 . We already discussed both the conditions for and interpretation of u_1 exhibiting quasi-concave differences in Sections 2.1 and 3.1. What is new is the condition in Lemma 5. Taking here condition 2. in Lemma 5, the intuition is straight-forward: if we think of x_2 as a measure of economic activity, condition 2 says that the MPC’s marginal payoff is increasing and concave in x_2 . Equivalently, the marginal loss is decreasing and convex in x_2 which means that higher activity lowers the marginal loss associated with raising the interest rate, but decreasingly so as the level of activity increases (the loss function is convex in x_2). This makes a lot of sense economically. All else equal, higher activity makes interest increases less painful in the eyes of an MPC, but this effect is bound to diminish as the economy progressively overheats.

Leaving a detailed “story” for future work, we may now add the public (agent $i = 2$), or we could add more than a single agent and extend (21) accordingly. We then ensure that the public’s objective function exhibits quasi-concave differences. For example, the public might be thought of as generating economic activity, and a concave policy function would mean that investments, etc., are decreasing at a steepening rate in the rate of interest. Adding the conditions of Lemma 5, we may then apply Theorem 3. This result tells us that if uncertainty increases (either for the central bank or the public) then in the Bayesian equilibrium both the interest rate and the activity variable decrease in the sense of second-order stochastic dominance. So if uncertainty increases for exogenous reasons, mean activity as well as the mean interest rate decrease, and both of these variables will fluctuate more.

4.2 The Bayesian Arms Race Model

In a Bayesian version of the classical arms race game from the field of conflict resolution (see *e.g.* Milgrom and Roberts (1990), p.1272), we ask here whether increased uncertainty about arms' effectiveness and opponents' intentions leads to an intensification of the arms race or not.

There are two countries, $i = 1, 2$, with identical state payoff functions $u_i(x_i, x_{-i}, z_i) = B(x_i - x_{-i} - z_i) - c x_i$. B is a strictly concave function and $c > 0$ a constant cost parameter. z_i is a random variable that reflects the relative effectiveness of the arms — real or imagined (for example a domestic media frenzy might correspond to a mean-preserving spread to z_i).³² Assuming that B is sufficiently smooth, we can use the conditions at the end of Theorem 3. By strict concavity, $D_{x_i} u(x_i, x_{-i}, z_i) = B'(x_i - x_{-i} + z_i) - c$ is increasing in x_{-i} , and the question is therefore whether it is also either convex or concave in (x_i, z_i) and in x_{-i} . Obviously, this depends entirely on whether B' is convex or concave, *i.e.*, on whether the third derivative of B is positive or negative. In the convex case (positive third derivative), the countries' policy functions are convex. Hence greater uncertainty will increase the affected country's (or countries') expected stock of arms as well as the variability (Lemma 4, which specifically says that given the other country's strategy, greater uncertainty will lead to a convex-increasing shift in the optimal distribution). The increased variability of the affected country's stock of arms will lead to a more uncertain environment for the other country and make *it* accumulate more arms (Lemma 5). This escalation continues until an equilibrium is reached with higher mean stocks of arms and greater uncertainty about the exact size of the arsenals (Theorem 3).³³ Note that a positive third derivative means that the countries are "prudent" (Kimball (1990)) — a well-understood behavioral trait that also plays a key role in other settings such as in income allocation problems (Carroll and Kimball (1996)). Of course, prudence, which in the words of Kimball (1990) (p.54) is "the propensity to prepare and forearm oneself in the face of uncertainty", has rather more beneficial consequences in income allocation models than it does in arms races. It is therefore not uniformly good news that experimental evidence seems to suggest that most people are prudent (Nussair et al (2011)). But of course, prudence may be situation-dependent or imprudent politicians may be elected. In this case B will have a negative third derivative, and the countries' policy functions will be concave so that greater uncertainty lowers the mean stock of arms in equilibrium. Note however, that according to Theorem 3, the variability will still increase, so whether decision makers are prudent or not, the risk of exceptionally high stocks of arms and the negative consequences in case of war still increases when the environment becomes more uncertain.

5 Stochastic Dynamic Programming

This section uses Theorem 2 to study the topic introduced in Section 2.2, namely the concavity/convexity of policy functions in stochastic dynamic programming problems. In income al-

³²A myriad of other specifications would of course be possible, for example costs could instead be random. This section's results may be applied for any such specification.

³³Note that since the conditions of Corollary 3 are satisfied, these conclusions are valid not just for mean-preserving spreads but for second-order stochastic dominance decreases more generally.

location problems, the relationship between earning risk and wealth accumulation is guided by whether the consumption function is concave or convex (Huggett (2004)). And in macroeconomics, it determines the effect of increased individual uncertainty on aggregate market outcomes in large dynamic economies (Acemoglu and Jensen (2015)) an example of which is the Aiyagari-model mentioned in Section 2.2 and returned to below.

The treatment and notation follows Chapter 9 of Stokey and Lucas (1989). The dynamic programming problem is,

$$(22) \quad \begin{aligned} \max \quad & \mathbb{E}_0[\sum_{t=0}^{\infty} \beta^t u(x_t, x_{t+1}, z_t)] \\ \text{s.t.} \quad & \begin{cases} x_{t+1} \in \Gamma(x_t, z_t), t = 0, 1, 2, \dots \\ \text{where } (x_0, z_0) \gg 0 \text{ are given.} \end{cases} \end{aligned}$$

In comparison with Stokey and Lucas (1989), two structural restrictions are made to simplify the exposition (but the results easily generalize, see Remarks 1-2 at the end of this section): first, the z_t 's are assumed to be *i.i.d.* with distribution μ_z ; second, only the one-dimensional case is considered, *i.e.*, it is assumed that $x_t \in X \subseteq \mathbb{R}$ and $z_t \in Z \subseteq \mathbb{R}$. We also maintain the standard assumption of a concave objective even though our general results in fact do not even require quasi-concavity (Remark 3 discusses this further).

Both X and Z are assumed to be convex sets equipped with their Borel σ -algebras.³⁴ The *value function* $v : X \times Z \rightarrow \mathbb{R}$ of (22) is,

$$(23) \quad v(x, z) = \sup_{y \in \Gamma(x, z)} \left[u(x, y, z) + \beta \int v(y, z') \mu_z(dz') \right].$$

The following assumptions are from Stokey and Lucas (1989), Chapter 9.

Assumption 2 $\Gamma : X \times Z \rightarrow 2^X$ is non-empty, compact-valued, continuous, and has a convex graph, *i.e.*, for all $x, \tilde{x} \in X, z \in Z$, and all $\lambda \in [0, 1]$: $\lambda y + (1 - \lambda)\tilde{y} \in \Gamma(\lambda x + (1 - \lambda)\tilde{x}, z)$ whenever $y \in \Gamma(x, z)$ and $\tilde{y} \in \Gamma(\tilde{x}, z)$.

Assumption 3 $u : X \times X \times Z \rightarrow \mathbb{R}$ is bounded and continuous, and $\beta \in (0, 1)$. Furthermore, $u(x, y, z)$ is concave in (x, y) and strictly concave in y .

Note that Assumption 2 in particular requires Γ to have a *convex graph*. As discussed in the first paragraph after Definition 3, this implies that Γ is a concave as well as a convex correspondence (*cf.* the conditions of Theorem 2). Under Assumptions 2-3, the value function $v = v(x, z)$ is uniquely determined, continuous, and concave in x . Furthermore, the *policy function* $g : X \times Z \rightarrow X$ is the well-defined and continuous function

$$(24) \quad g(x, z) = \arg \sup_{y \in \Gamma(x, z)} \left[u(x, y, z) + \beta \int v(y, z') \mu_z(dz') \right].$$

³⁴For our result on the policy function $g(x, z)$'s convexity in x , it may alternatively be assumed that Z is a countable set equipped with the σ -algebra consisting of all subsets of Z (see Stokey and Lucas (1989), Assumption 9.5.a.).

Theorem 4 (Convex Policy Functions in Dynamic Stochastic Programming Problems) Consider the stochastic dynamic programming problem (22) under Assumptions 2-3 and let $g : X \times Z \rightarrow X$ denote the policy function (24). Assume that $u(x, y, z)$ is differentiable and satisfies the following upper boundary condition: $\lim_{y \uparrow \sup \Gamma(x, z)} D_y u(x, y, z) = -\infty$ (or in some other way ensure that $\sup \Gamma(x, z)$ will never be optimal given (x, z)). Then the policy function g is convex in x if $D_x u(x, y, z)$ is non-decreasing in y and there exists a $k \geq 0$ such that $\frac{1}{1-k}[-D_y u(x, y, z)]^{1-k}$ is concave in (x, y) and $\frac{1}{1-k}[D_x u(x, y, z)]^{1-k}$ is convex in (x, y) .³⁵ If in addition $\Gamma(x, \cdot)$ is a convex correspondence and $\frac{1}{1-k}[-D_y u(x, y, z)]^{1-k}$ is concave in (y, z) , then the policy function g will also be convex in z .

Proof. See Appendix D. ■

Theorem 4 has a host of applications in macroeconomics. For example, it is applied in Acemoglu and Jensen (2015) to study how increased uncertainty affects the equilibria in large dynamic economies.

5.1 Income Allocation Models (Continued)

Recall the income allocation model from Section 2.2 whose Bellman equation corresponding to (23) is

$$(25) \quad v(x, z) = \max_{y \in \Gamma(x, z)} \tilde{u}((1+r)x + wz - y) + \beta \int v(y, z') \eta(dz').$$

Alternatively, we may express this in the form of (22),

$$(26) \quad \begin{aligned} \max \quad & \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \tilde{u}((1+r)x_t + wz_t - x_{t+1}) \right] \\ \text{s.t.} \quad & \begin{cases} x_{t+1} \in \Gamma(x_t, z_t), t = 0, 1, 2, \dots \\ \text{where } (x_0, z_0) \gg 0 \text{ are given.} \end{cases} \end{aligned}$$

Since x_t is savings at date t , the policy function $g((1+r)x + wz)$ is the *savings function* and the *consumption function* is

$$(27) \quad c((1+r)x + wz) = (1+r)x + wz - g((1+r)x + wz).$$

In terms of this section's general notation we have $u(x, y, z) = \tilde{u}((1+r)x + wz - y)$. It is easy to verify (and well known) that Assumptions 2-3 are satisfied. Under a standard boundary condition on \tilde{u} , we will never have $g((1+r)x + wz) = \sup \Gamma(x, z)$, *i.e.*, the consumer will not choose zero consumption at any date.

We have $D_y u(x, y, z) = -\tilde{u}'((1+r)x + wz - y)$ and $D_x u(x, y, z) = r \tilde{u}'((1+r)x + wz - y)$. $D_x u$ is strictly increasing in y since \tilde{u} is strictly concave. By Theorem 4, the savings function is convex (the consumption function is concave) if $\frac{1}{1-k}[\tilde{u}'((1+r)x + wz - y)]^{1-k}$ is concave and $\frac{1}{1-k}[r \tilde{u}'((1+r)x + wz - y)]^{1-k}$ is convex in (x, y) . When the Hessian determinants of $D_x u$ and $D_y u$ exist — *i.e.*, when \tilde{u} is thrice differentiable as in Carroll and Kimball (1996) — they both equal zero, and it is then straightforward to verify that the previous conditions on \tilde{u} will hold if and only if

$$(28) \quad \frac{\tilde{u}' \tilde{u}'''}{(\tilde{u}'')^2} = k \geq 0.$$

³⁵In the limit case $k = 1$, $\frac{1}{1-k}[f(x)]^{1-k}$ is by convention equal to $\log(f(x))$.

Thus we find the condition that \tilde{u} must be of the HARA-form introduced in Section 2.2. That \tilde{u} is of the HARA-form is also assumed by Carroll and Kimball (1996) who establish concavity of the consumption function when there are no borrowing limits. One easily verifies that the condition for convexity in z (last sentence of Theorem 4) is also satisfied when \tilde{u} is of the HARA-form. Thus we have generalized the result of Carroll and Kimball (1996) to the incomplete markets setting with borrowing constraints.

5.2 Inequality and the Lorenz Curve

We finish this section with a brief look at relationship between income inequality and concavity of consumption functions. Let $W_i = (1+r)x_i + wz_i$ denote income of agent i at a given moment in time and consider a continuum of agents $i \in [0, 1]$ with identical preferences but possibly different incomes. Each individual thus consumes $c(W_i)$. Letting η_W denote the frequency distribution of income, per-capita/mean consumption is then

$$\int c(W_i)\eta_W(dW_i).$$

From Atkinson (1970) we know that mean-preserving spreads to η_W are equivalent to increases in inequality in the sense of Lorenz dominance (Atkinson (1970), p. 246-247). When u is in the HARA class and the consumption function therefore concave, it follows from Jensen's inequality that a Lorenz increase in inequality implies *lower* per-capita consumption. But the Observations on page 5 allow us to go considerably further. Firstly, since the consumption function c is also increasing under standard conditions, the previous statement extends to generalized Lorenz dominance by Observation 3 on page 5.³⁶ Of more novelty, we can go beyond considerations of consumption per-capita. For example, by Observation 2, a Lorenz increase in inequality will lead to a second-order stochastic dominance decrease in the distribution of consumption. So we can conclude not only that the mean will decrease, but also that social welfare will decrease under any inequality adverse social welfare objective, *i.e.*, any social welfare objective that is increasing and concave in individual consumption levels (again see Atkinson (1970)). So we are able to study the relationship between the *inequality of opportunities* as embodied in Lorenz dominance and the *inequality of outcomes* as evaluated by a social welfare criterion.

6 Conclusion

This paper contributes to distributional comparative statics (DCS), *i.e.*, to the study of how changes in exogenous distributions affect endogenous distributions in economic models. Most DCS questions can be answered if suitable policy functions are either concave or convex. In the main theoretical contribution of the paper (Theorem 1), it is shown that concavity of the policy function hinges on an intuitive as well as easily verifiable condition on the primitives of a model, namely

³⁶The generalized Lorenz curve is constructed by scaling up the Lorenz curve by the distribution's mean and is equivalent to second-order stochastic dominance shifts, see *e.g.* Dorfman (1979).

quasi-concave differences. That observation parallels Topkis’ theorem (Topkis (1978)) which ensures that the policy function is increasing (strategic complementarity) when the objective function exhibits increasing differences. Theorem 2, as well as Theorem 5 in Appendix C, extends the result to policy correspondences (multiple optimizers) and multi-dimensional action sets.

Several areas of application were discussed including uncertainty comparative statics, international trade models of heterogenous firms (Melitz (2003)), the macroeconomic modeling of inequality, the issue of forecasting monetary policy, and stochastic dynamic programming. In all of these, the concavity of suitably defined policy functions turns out to drive the conclusions, ultimately owing to Observations 1-5 in Section 2.1 (page 5) and the fact that Lorenz dominance is equivalent to mean-preserving spreads and generalized Lorenz dominance is equivalent to decreases in second-order stochastic dominance (see the end of Section 5). As a concrete illustration of uncertainty comparative statics, a Bayesian arms race is studied and it is found that “prudence” (Kimball (1990)) determines whether mean stocks of arms increase or decrease when uncertainty goes up — but in all cases, a more uncertain environment also leads to greater uncertainty about the scale of destruction in the event of a war. The stochastic dynamic programming results are illustrated by generalizing a result due to Carroll and Kimball (1996) to allow for borrowing constraints. These results play a key role for distributional comparative statics in dynamic stochastic general equilibrium (DSGE) models — a theme taken up in Acemoglu and Jensen (2015) who study, for example, how increased uncertainty about future earnings prospects affects output per worker in the Aiyagari (1994) model.

7 Appendices

A Proof of Observations 1-5 in Section 2.1

The distribution μ_x is the image measure of μ under g . Hence $\int f(x)\mu_x(dx) = \int f(g(z))\mu(dz)$ for any function $f : X \rightarrow \mathbb{R}$ such that the integrals are well-defined. Each claim thus amounts to saying that for classes of functions \mathcal{F} and \mathcal{F}_x , if $\int h(z)\tilde{\mu}(dz) \geq \int h(z)\mu(dz)$ for all $h \in \mathcal{F}$, then $\int f(g(z))\tilde{\mu}(dz) \geq \int f(g(z))\mu(dz)$ for all $f \in \mathcal{F}_x$. In the case of Observation 1, \mathcal{F} and \mathcal{F}_x both equal the class of increasing functions and the claim follows from the fact that $f \circ g$ is increasing when both f and g are increasing. For Observation 2, \mathcal{F} is the class of convex functions and \mathcal{F}_x the class of decreasing, convex functions, and the claim follows because $f \circ g$ is convex when g is concave and f is convex and decreasing. For Observation 3, both \mathcal{F} and \mathcal{F}_x equal the class of increasing, concave functions and the conclusion follows because $f \circ g$ is increasing and concave when both f and g are increasing and concave. Observations 4-5 are proved by the same arguments as Observations 2-3 and may be omitted.

B Proof of Lemma 1

To facilitate the results on multi-dimensional strategy sets in Appendix C, the general case where $X \subseteq \mathbb{R}^n$ is considered. We thus prove that if u is differentiable in x , it will exhibit quasi-concave

differences if and only if $D_{x_j} u(x, z)$ is quasi-concave in (x, z) for all $j = 1, \dots, n$ (evidently, Lemma 1 is a special case of this statement).

“ \Rightarrow ”: Since $D_{x_j} u(x, z) = \lim_{\delta \rightarrow 0} \frac{u(x + \delta \epsilon_j, z) - u(x, z)}{\delta}$ where ϵ_j denotes the j 'th unit vector, and quasi-convexity is preserved under pointwise limits (Johansen (1972)), each partial derivative $D_{x_j} u(x, z)$ is quasi-convex at (x, z) when u exhibits quasi-convex differences at (x, z) . “ \Leftarrow ”: This direction is not easy. The idea is to prove the contrapositive by contradiction (note that since quasi-convexity is not preserved under integration, we cannot use the fundamental theorem of calculus). So we assume that u does not exhibit quasi-convex differences, that each partial derivative $D_{x_j} u(x, z)$ is quasi-convex, and then derive a contradiction. For $\alpha \in [0, 1]$ set $x_\alpha \equiv \alpha x_0 + (1 - \alpha)x_1$ and $z_\alpha = \alpha z_0 + (1 - \alpha)z_1$. Say that u exhibits quasi-convex differences in the direction $\eta > 0$ at (x_0, z_0, α) if for all $\delta_n > 0$ in some neighborhood of 0:

$$(29) \quad u(x_\alpha + \delta_n \eta, z_\alpha) - u(x_\alpha, z_\alpha) \leq \max\{u(x_0 + \delta_n \eta, z_0) - u(x_0, z_0), u(x_1 + \delta_n \eta, z_1) - u(x_1, z_1)\}$$

It is easy to see that if u exhibits quasi-convex differences (on all of $X \times Z$), then it exhibits quasi-convex differences in all directions $\eta > 0$ at all $(x, z, \alpha) \in X \times Y \times [0, 1]$. Let ϵ_j denote the j 'th unit vector (a vector with 1 in the j 'th coordinate and zeroes everywhere else). Since a function is quasi-convex in all directions if and only if it is quasi-convex in all unit/coordinate directions ϵ_j , we may (as always) restrict attention to the directions of the coordinates in the previous statement. Hence if u does not exhibit quasi-convex differences, there will exist a coordinate direction ϵ_j , $(x_0, z_0), (x_1, z_1) \in X \times Y$, $\hat{\alpha} \in [0, 1]$ and a sequence $\delta_n \downarrow 0$ such that for all n :

$$(30) \quad u(x_{\hat{\alpha}} + \delta_n \epsilon_j, z_{\hat{\alpha}}) - u(x_{\hat{\alpha}}, z_{\hat{\alpha}}) > \max\{u(x_0 + \delta_n \epsilon_j, z_0) - u(x_0, z_0), u(x_1 + \delta_n \epsilon_j, z_1) - u(x_1, z_1)\}$$

Note that we necessarily have $\hat{\alpha} \in (0, 1)$ when the previous inequality holds. Intuitively, the inequality says that there exists a point (x_α, z_α) on the line segment between (x_0, z_0) and (x_1, z_1) at which $u(\cdot + \delta_n \epsilon_j, \cdot) - u(\cdot, \cdot)$ takes a strictly higher value than at any of the endpoints. Now, divide through (30) with δ_n and take limits:

$$D_{x_j} u(x_{\hat{\alpha}}, z_{\hat{\alpha}}) \geq \max\{D_{x_j} u(x_0, z_0), D_{x_j} u(x_1, z_1)\}$$

Since $D_{x_j} u(\cdot, \cdot)$ is quasi-convex, it follows that: $D_{x_j} u(x_{\hat{\alpha}}, z_{\hat{\alpha}}) = \max\{D_{x_j} u(x_0, z_0), D_{x_j} u(x_1, z_1)\}$. Assume without loss of generality that $D_{x_j} u(x_0, z_0) \geq D_{x_j} u(x_1, z_1)$. Since $D_{x_j} u(x_0, z_0)$ is quasi-convex and $0 < \hat{\alpha} < 1$, it follows that either (i) $D_{x_j} u(x_0, z_0) = D_{x_j} u(x_\alpha, z_\alpha)$ for all $\alpha \in [0, \hat{\alpha}]$ or (ii) $D_{x_j} u(x_1, z_1) = D_{x_j} u(x_\alpha, z_\alpha)$ for all $\alpha \in [\hat{\alpha}, 1]$ (or both).³⁷ Consider case (i) (the proof in case (ii) is similar). When (i) holds, u 's restriction to the line segment between (x_0, z_0) and $(x_{\hat{\alpha}}, z_{\hat{\alpha}})$ must necessarily be of the form: $u(x, z) = c x_j + g(x_{-j}, z)$ where $c_n = D_{x_n} u(x_0, z_0)$ (a constant) and x_{-j}

³⁷A quasi-convex function's restriction to a convex segment as the one considered here can always be split into two segments, one which is non-increasing and one which is non-decreasing (and in the present situation, there *must* first be non-increasing segment since the function's value weakly decreases between the endpoints). On the convex line segment between (x_0, z_0) and (x_1, z_1) we have in the present situation that the function begins at $D_{x_j} u(x_0, z_0)$, again takes the value $D_{x_j} u(x_0, z_0)$ at (x_α, z_α) and then moves to a weakly lower value $D_{x_j} u(x_1, z_1)$ at the end-point (x_1, z_1) . It follows that if $D_{x_j} u(\cdot, \cdot)$ is not constant on the first interval (corresponding to $\alpha \in [0, \hat{\alpha}]$) it must strictly decrease and then strictly increase on this interval, which implies that $D_{x_j} u(\cdot, \cdot)$ is constant on the second of the two intervals.

denotes all coordinates of x except for the j 'th one (remember that j is the positive coordinate of e_j). But then $u(x_{\hat{\alpha}} + \delta_n \epsilon_j, z_{\hat{\alpha}}) - u(x_{\hat{\alpha}}, z_{\hat{\alpha}}) = u(x_0 + \delta_n \epsilon_j, z_0) - u(x_0, z_0) = c \delta_n \epsilon_j$ which implies that:

$$(31) \quad u(x_{\hat{\alpha}} + \delta_n \epsilon_j, z_{\hat{\alpha}}) - u(x_{\hat{\alpha}}, z_{\hat{\alpha}}) \leq \max\{u(x_0 + \delta_n \epsilon_j, z_0) - u(x_0, z_0), u(x_1 + \delta_n \epsilon_j, z_1) - u(x_1, z_1)\}$$

Comparing (31) with (30) we have a contradiction, and the proof is complete.

C Multi-dimensional Decision Variables

This appendix treats concavity of the policy correspondence in the case where the decision vector is allowed to live in an arbitrary ordered topological vector lattice V , $x \in X \subseteq V$. Note that in this setting, the order mentioned in Definition 3 is the order inherited from V . So if $V = \mathbb{R}^N$ with the usual Euclidean order, the theorem below implies that each coordinate correspondence $G_n : Z \rightarrow 2^{X^n} \subseteq 2^{\mathbb{R}}$ is concave in the sense discussed in detail in Section 3.2.

The multi-dimensional setting forces us to make some additional assumptions. In comparison with Theorem 2, X must be a lattice, u must be supermodular in the decision vector, and Γ must be upper semi-lattice valued as well as order convex valued rather than merely convex valued.³⁸ Finally, the boundary conditions must be suitably generalized as we turn to first. It should be noted that all of these assumptions automatically are satisfied when X is one-dimensional. Hence the result to follow encompasses Theorem 2.

First, the basic definitions. Say that a point $x \in \Gamma(z)$ lies on the *upper* [*lower*] *boundary* of $\Gamma(z)$ if there does not exist an $x' \in \Gamma(z)$ with $x' \gg x$ ($x' \ll x$). The upper boundary is denoted by $\bar{\mathbf{B}}(\Gamma(z))$ and the lower boundary is denoted by $\underline{\mathbf{B}}(\Gamma(z))$. What we are going to require in the theorem below is precisely as in Theorem 2 except that the infimum is replaced with the lower boundary.

Next, X must be a *lattice*, *i.e.*, if x and x' lie in X so do their infimum $x \wedge x'$ and supremum $x \vee x'$. If $X \subseteq \mathbb{R}^n$ with the usual Euclidean/coordinatewise order, the infimum (supremum) is simply the coordinatewise minimum (maximum). Assuming that X is a lattice is actually a very weak additional requirement in the present framework because it is the constraint correspondence Γ that determines the feasible set. It is the next assumption that really has “bite”. A lower semi-lattice [upper semi-lattice] is a subset $A \subset X$ with the property that if $x, x' \in A$ then the infimum $x \wedge x'$ [supremum $x \vee x'$] also lies in A . Either is of course weaker than being a lattice. The lower or upper semi-lattice is *order-convex* if $a, b \in A$, and $a \leq a' \leq b$ imply $a' \in A$ (in words, if the set contains an ordered pair of elements, it contains the entire order interval between these elements). Note that order-convexity is stronger than convexity in general, although the two coincide in the one-dimensional case. A budget set is an order-convex lower semi-lattice (it is not a lattice), and a firm's input requirement set is an order-convex upper semi-lattice (but again not a lattice). As these examples indicate, the fact that we avoid assuming that Γ 's values are sublattices of X greatly expands the scope of the theorem below.

Finally, u must be supermodular in the decision variables. The well-known definition is as follows.

³⁸Note that this once again precisely parallels monotone comparative statics. In that setting supermodularity and lattice-type assumptions are also unnecessary/trivially satisfied in the one-dimensional case but must be imposed in multiple dimensions.

Definition 4 (Topkis (1978)) *The objective function $u : X \times Z \rightarrow \mathbb{R}$ is supermodular in x if $u(x \vee x', z) + u(x \wedge x', z) \geq u(x, z) + u(x', z)$ for all $x, x' \in X$ and for all $z \in Z$. If u is twice differentiable in x and $X \subseteq \mathbb{R}^n$, it is supermodular in x if and only if the Hessian matrix $D_{xx}^2 u(x, z) \in \mathbb{R}^{n \times n}$ has non-negative off-diagonal elements (for all x and z).*

We are now ready to state and prove the main result with multi-dimensional action sets. Note that as in Theorems 1-2, the boundary condition is trivially satisfied if the optimization problem is unrestricted or attention is restricted to interior solutions.

Theorem 5 (Concavity of the Policy Correspondence, Multidimensional Case) Let Z be a convex subset of a real vector space and X a convex lattice. Define the policy correspondence $G(z) = \operatorname{argsup}_{x \in \Gamma(z)} u(x, z)$ where $u : X \times Z \rightarrow \mathbb{R}$ is supermodular in x and $\Gamma : Z \rightarrow 2^X$ has order-convex values. Assume that $G(z)$ is non-empty and compact for all $z \in Z$. Then:

1. The policy correspondence G is concave if $u : X \times Z \rightarrow \mathbb{R}$ exhibits quasi-concave differences, Γ is concave and upper semi-lattice valued, and $x \in G(z) \Rightarrow x \notin \underline{B}(\Gamma(z))$ for all $z \in Z$.
2. The policy correspondence G is convex if $u : X \times Z \rightarrow \mathbb{R}$ exhibits quasi-convex differences, Γ is convex and lower semi-lattice valued, and $x \in G(z) \Rightarrow x \notin \overline{B}(\Gamma(z))$ for all $z \in Z$.

Proof of Theorem 5. The convex case 2. is proved (the proof of the concave case is similar). Pick $z_1, z_2 \in Z$, $x_1 \in G(z_1)$, and $x_2 \in G(z_2)$. Exactly as in the proof of Theorem 1, we can use quasi-convex differences to conclude that for some $\delta \gg 0$, $u(x_\alpha, z_\alpha) \geq u(x_\alpha + \delta, z_\alpha)$ for all $\alpha \in [0, 1]$. Hence by quasi-concavity of u in x , $u(x, z_\alpha)$ is non-increasing for $x \geq x_\alpha$. We wish to show that for all α there exists $\hat{x} \in G(z_\alpha)$ with $\hat{x} \leq x_\alpha$. Pick any $x \in G(z_\alpha)$. It is first shown that

$$(32) \quad x \wedge x_\alpha \in \Gamma(z_\alpha).$$

Since Γ has convex values, there exists some $\tilde{x} \in \Gamma(z_\alpha)$ with $\tilde{x} \leq x_\alpha$. We have $x \in \Gamma(z_\alpha)$ (since $x \in G(z_\alpha)$) and so since Γ 's values are lower semi-lattices, $x \wedge \tilde{x} \in \Gamma(z_\alpha)$. But $x \wedge \tilde{x} \leq x \wedge x_\alpha \leq x$, hence $x \wedge x_\alpha \in \Gamma(z_\alpha)$ because Γ has order-convex values. That was what we wanted to show. Next use supermodularity of $u(\cdot, t)$ and the fact that $u(\cdot, z_\alpha)$ is non-increasing for $x \geq x_\alpha$ (implies that $u(x_\alpha, z_\alpha) \geq u(x \vee x_\alpha, z_\alpha)$) to conclude that:

$$(33) \quad u(x, z_\alpha) - u(x \wedge x_\alpha, z_\alpha) \leq u(x \vee x_\alpha, z_\alpha) - u(x_\alpha, z_\alpha) \leq 0$$

(32)-(33) imply that $x \wedge x_\alpha \in G(z_\alpha)$. But since clearly $x \wedge x_\alpha \leq x_\alpha$ this completes the proof. ■

D Proof of Theorem 4

The value and policy functions will equal the pointwise limits of the sequences $(v^n)_{n=0}^\infty$ and $(g^n)_{n=0}^\infty$ determined by:

$$(34) \quad v^{n+1}(x, z) = \sup_{y \in \Gamma(x, z)} \left[u(x, y, z) + \beta \int v^n(y, z') \mu_z(dz') \right], \text{ and}$$

$$(35) \quad g^n(x, z) = \arg \sup_{y \in \Gamma(x, z)} \left[u(x, y, z) + \beta \int v^n(y, z') \mu_z(dz') \right].$$

Under the Theorem's conditions, v^n is concave in x for all n . Since a concave function is absolutely continuous, v^n will be absolutely continuous in x for all n . The following result presents conditions under which an absolutely continuous function is quasi-concave. It generalizes the sufficiency part of Lemma 1.

Lemma 6 Assume that $u : X \times Z \rightarrow \mathbb{R}$ is absolutely continuous in $x \in X$, *i.e.*, assume that $u(x, z) = \alpha(a, z) + \int_a^x p(\tau, z) d\tau$ for a Lebesgue integrable function $p : X \times Z \rightarrow \mathbb{R}$ (here $a = \inf X$ and α is a function that does not depend on x). Then u exhibits quasi-convex differences [quasi-concave differences] if $p(x, z)$ is quasi-convex [quasi-concave].

Proof. The statement can be verified by going through the proof of Lemma 1 and everywhere replace u 's derivative with p . ■

Next we need a result on how to ensure that the sum of two functions exhibits quasi-concave or quasi-convex differences.

Lemma 7 Let u be of the form $u(x, z) = f(x, z) + h(-x, z)$ where $f, h : X \times Z \rightarrow \mathbb{R}$ are differentiable with $D_x f(x, z) = f'_x(x, z) \geq 0$ and $D_x h(x, z) = -h'_x(-x, z) \leq 0$. Then u exhibits quasi-convex differences on the subset $\{(x, z) \in X \times Z : D_x u(x, z) \leq 0\}$ if there exists a $k \geq 0$ such that $\frac{1}{1-k}[f'_x]^{1-k}$ is convex and $\frac{1}{1-k}[h'_x]^{1-k}$ is concave.³⁹ If instead $\frac{1}{1-k}[f'_x]^{1-k}$ is concave and $\frac{1}{1-k}[h'_x]^{1-k}$ is convex, then u exhibits quasi-concave differences on the subset $\{(x, z) \in X \times Z : D_x u(x, z) \geq 0\}$.

Proof. In the online appendix (Jensen (2015), Lemma 3). ■

From now on, call a function $f : X \rightarrow \mathbb{R}$ *k-convex* [*k-concave*] if $\frac{1}{1-k}[f(x)]^{1-k}$ is convex [concave] where, as previously mentioned, the case $k = 1$ is taken to mean log-convex [log-concave] by convention.⁴⁰

Lemma 8 Given z , let a denote the least point at which $v^n(\cdot, z)$ is defined. Assume that $-D_y u(x, y, z)$ is k -concave in (x, y) [k -concave in (y, z)], and that $v^n(y, z) = \int_a^y p^n(\tau, z) d\tau$ where $p^n(\cdot, z)$ is k -convex. Then $g^n(x, z)$ is convex in x [convex in z].

Proof. This is a direct application Theorem 2's corollary (Corollary 2) to the optimization problem in (35). We consider here only the convexity of g^n in x (the exact same argument implies convexity in z under the Lemma's square-bracketed assumption). Except for quasi-convex differences in (x, y) , all the assumptions of Corollary 2 are clearly satisfied (in particular, Γ is a convex correspondence as mentioned immediately after Assumption 3). To see that quasi-convex differences

³⁹In the limit case $k = 1$, $\frac{1}{1-k}[f'_x]^{1-k}$ is by convention equal to $\log(f'_x)$ (similarly for h'_x).

⁴⁰Quite a bit can be said about such functions, but since this is mainly a mathematical distraction from the point of view of this paper, further investigation has been relegated to the online appendix.

holds, we use Lemma 6 and must thus verify that,

$$(36) \quad D_y u(x, y, z) + \beta \int p^n(y, z') \mu_z(dz'),$$

is quasi-convex in (x, y) on the relevant set which, allowing for solutions at lower boundary points is $A \equiv \{(x, y) \in X^2 : D_y u(x, y, z) + \beta \int D_y v^n(y, z') \mu_z(dz') \leq 0\}$. To see that this holds, first use that k -convexity is preserved under integration (Jensen (2015), Lemma 2) to conclude that when $p^n(y, z')$ is k -convex in y , $\beta \int p^n(y, z') \mu_z(dz')$ is k -convex in y . Then use Lemma 7. ■

To finish the proof we need just one last technical result.

Lemma 9 Assume that $D_x u(x, y, z)$ is k -convex in (x, y) and non-decreasing in y and that $g^n(x, z)$ is convex in x . Then $v^{n+1}(x, z) = \int_a^x p^{n+1}(\tau, z) d\tau$ where $p^{n+1}(\cdot, z)$ is k -convex.

Proof. Since v^{n+1} is absolutely continuous, we can (abusing notation slightly) write it as: $v^{n+1}(x, z) = \int_a^x D_x v^{n+1}(\tau, z) d\tau$. In particular, $D_x v^{n+1}(x, z)$ exists almost everywhere and when it exists $D_x v^{n+1}(x, z) = D_x u(x, g^n(x, z), z)$ by the envelope theorem. k -convexity of $p^{n+1}(x, z) \equiv D_x u(x, g^n(x, z), z)$ in x now follows immediately from the fact that k -convexity is preserved under convex, increasing transformations (Jensen (2015), Lemma 1). ■

To prove that g is convex, consider the value and policy function iterations (34)-(35). Start with any value function v^0 such that $v^0(y, z) = \int_a^y p^0(\tau, z) d\tau$ where $p^0(\cdot, z)$ is k -convex. Then by Lemma 8, g^0 is convex. Hence by Lemma 9, $v^1(y, z) = \int_a^y p^1(\tau, z) d\tau$ where $p^1(\cdot, z)$ is k -convex. Repeating the argument, g^1 is convex and $v^2(y, z) = \int_a^y p^2(\tau, z) d\tau$ where $p^2(\cdot, z)$ is k -convex. And so on ad infinitum. The pointwise limit of a sequence of convex function is convex, hence $g(\cdot, z) = \lim_{n \rightarrow \infty} g^n(\cdot, z)$ is convex. The same argument applies for convexity in z , concluding the proof of Theorem 4.

Remark 1 (General Markov Processes) The previous proof goes through without any modifications if z_t is allowed to be a general Markov process, *i.e.*, if the functional equation (23) is replaced with:

$$(37) \quad v(x, z) = \sup_{y \in \Gamma(x, z)} [u(x, y, z) + \beta \int v(y, z') Q(z, dz')],$$

where Q is z_t 's transition function. Indeed, the previous proof goes through line-by-line if we instead begin with the functional equation (37).

Remark 2 (Multidimensional Strategy Sets) The proof also easily extends to the case where X and Z are multidimensional (a case treated in the Appendix). The only modification needed is in the proof of Lemma 8 where now Theorem 5 in the appendix is needed to conclude that g^n is convex, in place of Theorem 2. Thus Theorem 4 extends to the multidimensional case if we in addition assume that u is supermodular in y , that Γ 's values are lower semi-lattices, and that optimizers stay away from the upper boundary.⁴¹

⁴¹In particular, the objective function in (35) is supermodular in y when u is supermodular in y because supermodularity/increasing differences is preserved under integration (Topkis (1998), Theorem 2.7.6.) and $v^n(y, z')$ is therefore supermodular in y for all n .

Remark 3 (Dispensing with Concavity Assumptions) The results of this section can be extended to cases where the objective function is not concave in (x, y) and/or is not strictly concave in y (see Assumption 3). What is critical for the proof is that the value function is absolutely continuous so that Lemma 6 can be applied. Of course, absolute continuity of the value function can be established under weak direct conditions (see *e.g.* Milgrom and Segal (2002)) instead of being derived as a consequence of concavity as done above. Instead of a convex policy function one then gets a convex policy correspondence.

References

- Acemoglu, D. and M.K. Jensen (2013): “Aggregate Comparative Statics”, *Games and Economic Behavior* 81, 27-49. 6
- Acemoglu, D. and M. K. Jensen (2015): “Robust Comparative Statics in Large Dynamic Economies”, *Journal of Political Economy* 123, 587-640. 2, 3, 6, 7, 21, 24, 25, 27
- Aiyagari, S.R. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving”, *Quarterly Journal of Economics* 109, 659-684. 0, 1, 2, 7, 27
- Athey, S. (2002): “Monotone Comparative Statics under Uncertainty”, *Quarterly Journal of Economics* 117, 187-223. 19
- Atkinson, A.B. (1970): “On the Measurement of Inequality”, *Journal of Economic Theory* 2, 244-263. 1, 26
- Carroll, C.D. and M.S. Kimball (1996): “On the Concavity of the Consumption Function”, *Econometrica* 64, 981-992. 0, 2, 7, 23, 25, 26, 27
- Carroll, C.D. and M.S. Kimball (2001): “Liquidity Constraints and Precautionary Saving”, *Working Paper*. 7
- Dorfman, R. (1979): “A Formula for the Gini Coefficient”, *Review of Economics and Statistics* 61, 146-49. 26
- Heijdra, B.J., *Foundations of Modern Macroeconomics*, Oxford University Press, 2nd Edition, 2009. 3
- Huggett, M. (2004): “Precautionary Wealth Accumulation”, *Review of Economic Studies* 71, 769-781. 3, 7, 24
- Jensen, M.K. (2012): “Distributional Comparative Statics”, University of Birmingham Discussion paper # 12-08. 12
- Jensen, M.K. (2015): “The Class of k -Convex Functions”, *Online Supplement*. 31, 32
- Johansen, S. (1972): “A Representation Theorem For a Convex Cone of Quasi Convex Functions”, *Math. Scand.* 30, 297-312. 28
- Kimball, M.S. (1989): “The Effect of Demand Uncertainty on a Precommitted Monopoly Price”, *Economics Letters* 30, 1-5.
- Kimball, M.S. (1990): “Precautionary Saving in the Small and in the Large”, *Econometrica* 58, 53-73. 23, 27

- Kuroiwa, D. (1996): "Convexity of Set-Valued Maps", *Applied Mathematics Letters* 9, 97-101. 14
- Kydland, F. and E.C. Prescott (1977): "Rules Rather than Discretion: The Inconsistency of Optimal Plans", *Journal of Political Economy* 85, 473-492. 1, 3
- Leahy, J.V. and T.M. Whited (1996): "The Effect of Uncertainty on Investment: Some Stylized Facts", *Journal of Money, Credit, and Banking* 28, 64-83. 17
- Mas-Collel, A., M.D. Whinston, and J.R. Green, *Microeconomic Theory*, Oxford University Press, New York, 1995. 16
- Melitz, M.J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity", *Econometrica* 71, 1695-1725. 0, 1, 2, 17, 27
- Milgrom, P. and J. Roberts (1990): "Rationalizability, and Learning, and Equilibrium in Games with Strategic Complementarities", *Econometrica* 58, 1255-1278. 23
- Milgrom, P. and J. Roberts (1994): "Comparing Equilibria", *American Economic Review* 84, 441-459. 2, 6
- Milgrom, P. and I. Segal (2002): "Envelope Theorems for Arbitrary Choice Sets", *Econometrica* 70, 583-601. 33
- Milgrom, P. and C. Shannon (1994): "Monotone Comparative Statics", *Econometrica* 62, 157-180. 2, 5, 6, 16
- Myatt, D.P. and C. Wallace (2014): "Central Bank Communication Design in a Lucas-Phelps Economy", *Journal of Monetary Economics* 63, 64-79. 4
- Noussair, C.N., S.T. Trautmann, and Gijs v.-d. Kuilen (2011): "Higher Order Risk Attitudes, Demographics, and Financial Decisions", Working Paper, University of Tilburg. 23
- Quah, J. (2007): "The Comparative Statics of Constrained Optimization Problems", *Econometrica* 75, 401-431. 2
- Quah, J. and B. Strulovici (2012): "Aggregating the Single Crossing Property", *Econometrica* 2333-2348. 9
- Rothschild, M. and J.E. Stiglitz (1971): "Increasing Risk II: Its Economic Consequences", *Journal of Economic Theory* 3, 66-84. 19
- Shaked, M. and J.G. Shanthikumar, *Stochastic Orders*, Springer, New York, 2007. 4
- Smithson, R.E. (1971): "Fixed Points of Order Preserving Multifunctions", *Proceedings of the American Mathematical Society* 28, 304-310. 20
- Stokey, N.E. and R.E. Lucas jr. (with E.C. Prescott), *Recursive Methods in Economic Dynamics*, Harvard University Press, Cambridge, 1989. 8, 24
- Suen, Richard (2015): "Concave Consumption Function and Precautionary Wealth Accumulation", *Working Paper*, University of Leicester. 7
- Topkis, D.M. (1978): "Minimizing a Submodular Function on a Lattice", *Operations Research*, 26, 305-321. 2, 5, 8, 14, 16, 20, 27, 30
- Topkis, D.M., *Supermodularity and Complementarity*, Princeton University Press, New Jersey, 1998. 32

Van Zandt, T. and X. Vives (2007): "Monotone Equilibria in Bayesian Games of Strategic Complementarities" *Journal of Economic Theory* 134, 339-360. 5, 21

Varian, H.R., *Microeconomic Analysis, 3rd Edition*, New York: W.W. Norton, 1992. 10