Journal of Mathematical Economics 18 (1989) 303–324. North-Holland

DECENTRALIZED ALLOCATION OF RESOURCES AMONG MANY PRODUCERS

Zvi ARTSTEIN

The Weizmann Institute of Science, Rehovot 76100, Israel

Roger J-B WETS

University of California, Davis, CA 95616, USA

Submitted September 1987, accepted November 1988

We reconsider the question of allocating resources in large teams using decentralized procedures. Our assumption on the distribution of producers is that it approximates a given underlying distribution in producers space; thus we relax prior approaches which utilize replicas or iid samplings. We examine when solving the allocation problem for the underlying distribution yields an appropriate solution to the specific sample. We then show how market-like mechanisms may be used to get a decentralized decision process which is asymptotically optimal.

1. Introduction

We examine in this paper the possibility of using market-like structures in allocation problems with a large number of producers. The model follows Arrow and Radner (1979) and Groves and Hart (1982); we now describe the motivation.

While allocating resources in an optimal way is desired, it may result with a massive amount of communication if the number of participants is large. A market-like structure may help. For instance, prices may be announced, according to which the agents adjust their demands and local decisions. Under natural conditions such a price mechanism can indeed generate the optimal allocation. However, the computation of these prices may require an effort, e.g., collecting data from the many agents, a burden that could be beyond our means and better be eliminated.

A situation where this drawback may be overcome is where there is a good idea about the statistical distribution of the agents' characteristics. Intuitively at least, if the agents form a good sample of a known population, then the parameters of the latter may be used in the allocation process, resulting in an approximation to the optimal solution. The, hopefully minor, loss in production is then compensated by the simplicity of the process.

This idea was pursued by Arrow and Radner (1979) in a framework where resources ought to be allocated among a large number of producers, each of them is subject to local decisions; the objective being the maximization of the total output. The sampling property mentioned above is obtained in Arrow and Radner (1979) by letting the producers be independently drawn from a given population. The result is that, almost surely, the averaged optimal value obtained with full exchange of information coincides in the limit with the averaged optimal value when each of the producers makes his/her own decision. The common value in the limit is the one obtained by considering an optimal allocation for the underlying population with the given density. Groves and Hart (1982) complemented the result by providing asymptotically optimal allocation schemes. The schemes work, almost surely, as it is assumed [like in Arrow and Radner (1979)] that the producers are identically distributed and randomly drawn from a given distribution; another assumption was that the underlying distribution is known to each of the producers.

In this paper we try to relax the latter two assumptions. We replace the probabilistic framework by an assumption that the set of producers forms a good sample, how this is done is explained in section 3. This eliminates the assumptions of independence and identical distribution, and leads to deterministic, rather than almost sure, conclusions.

The improvement is more than a mere technicality; it has a transparent economic meaning. Samplings are fundamental to the type of models that we analyze. In our model the sampling occurs when producers come to pick their allocated resources. Such procedures occur in real life. For instance, a government may announce a loan policy for a developing area (policy which may be modified in time). Prospective developers, from the space of developers, come then to get the loans. There is no reason why the developers that show up would form an identically distributed and independent sampling. An assumption concerning the probability distribution is more reasonable. (Even in a statistical framework, when a sample is picked out of a large population, the validity of the iid assumption in real practice is always in doubt.)

Apparently, Arrow and Radner (1979) were bothered by the strong probabilistic assumptions, and gave an example showing that the iid assumption cannot be dropped. In this paper we managed to get by without the iid assumption by introducing another condition, considerably relaxed, which reflects an economic characteristic of the market. We call it p-tightness of the sequence, where p is the vector of shadow prices for the allocation process. The condition is introduced and explained in section 5.

A variation of the model allows the optimal policy to be modified as the producers come along. This is important in particular when information on the distribution of agents is not known beforehand, and has to be accumulated during the execution of the allocation – a real life situation. When modifying the allocation schemes of Groves and Hart (1982) we develop also such adaptive processes.

We organize the paper as follows. In the next section we introduce the basic model, following Arrow and Radner (1979) and Groves and Hart (1982), and give the technical assumptions. Our approach to the problem is displayed in section 3, and compared with the approach of Arrow and Radner. In section 4 an example is constructed, showing that a useful property does not hold without the aforementioned additional condition. The latter has an economic interpretation, and we assume it in the rest of the paper. In section 5 we establish a convergence result. In section 6 we consider allocation schemes in our relaxed framework, following and modifying the schemes suggested in Groves and Hart (1982). The closing section displays the adaptive version of the allocation scheme.

2. The model

We first set the model, along the lines of Arrow and Radner (1979) and Groves and Hart (1982), and describe the optimization problems that arise. The technical assumptions follow.

Producers $t_1,\dots,t_n$ come from a set T of agents. Each producer, say t, when participating in the allocation process, has to choose an element ℓ from a given set $L(t)$ of decisions; these are called local decisions and reflect, say, labor recruits, facility installments, etc. Then, if provided with the vector x of resources, the agent t produces the amount $F(x,\ell,t)$. To simplify notations we allow ℓ to belong to a fixed set L [which contains the union of all $L(t)$] and set the production function $F(x,\ell,t)=-\infty$ if $\ell\notin L(t)$. The vector x of resources belongs to a set $X\subset R^s$, the s-dimensional euclidean space. Thus

$$F(x,\ell,t)\colon X\times L\times T\to[-\infty,\infty).$$

We wish to examine the case where the total amount of resources to be allocated in the process is large for a large number of producers. Since we plan to consider the case where $n\to\infty$, it is convenient to express the products in terms of averages. We denote by

$$q^0=(q_1,\dots,q_s)$$

the vector of average of quantities to be allocated.

Let $N=\{t_1,\dots,t_n\}$ be a set of producers. (Here t_i denotes both the ith producer and the parameter in T assigned to the ith producer. We leave this formal ambiguity; no confusion in the interpretation of the results should arise.) The objective of the allocation procedure is to maximize the total output, or equivalently, the average output. We assume that $F(x,\ell,t)$ is non-decreasing in x, which reflects free disposal; therefore the optimization problem facing the ensemble N is

$$P(N)\begin{cases}\text{maximize } \frac{1}{n}\sum_{i=1}^{n} F(x_i,\ell_i,t_i)\\ \text{subject to } \frac{1}{n}\sum_{i=1}^{n} x_i = q^0 \text{ and } \ell_i \in L.\end{cases}$$

We denote by $v(N)$ the supremum of the problem $P(N)$.

A centrally monitored allocation with full control over the choices ℓ_i and x_i can achieve the value $v(N)$, or get output as close to $v(N)$ as desired if $P(N)$ does not have an optimal solution. As explained in the introduction, we wish to come up with a value close to $v(N)$ using a process which does not require centrally monitored decisions. We examine a case where T, the set of potential producers, is equipped with a probability distribution (and the sample $t_1,\dots,t_n$ will eventually be related to this probability distribution). Given a probability distribution μ on T, we consider the optimization problem

$$P(\mu)\begin{cases}\text{maximize } \int_T F(x(t),\ell(t),t)\,\mathrm{d}\mu(t)\\ \text{subject to } \int_T x(t)\,\mathrm{d}\mu(t) = q^0.\end{cases}$$

We denote by $v(\mu)$ the superemum of $P(\mu)$. [Needless to say, the functions $(x(t),\ell(t))$ participating in the problem are assumed measurable, and so that $x(t)$ and $F(x(t),\ell(t),t)$ are integrable. Such pairs are called *admissible policies.*]

The previous setting is for any probability measure on T; we denote by μ^* the specific distribution that T originally is equipped with. Trying to follow the motivation given in the introduction, we face two questions. First:

Is $v(\mu^*)$ a good approximation of $v(N)$ for n large?

If the answer is positive we ask:

How can a solution of $P(\mu^*)$ be employed in simplifying the allocation process for N, yielding a good approximation of $v(N)$?

The analysis of the first query is done in the next three sections, and that of the second question in the final two sections. Here we display the technical conditions.

Assumptions

(i) T is a complete separable metric space, considered with its Borel structure.
(ii) L is a complete separable metric space.
(iii) X is closed convex subset of R^s, bounded from below.
(iv) q^0 is in the interior of X.

In Arrow and Radner (1979) and Groves and Hart (1982), $X = R^s_+$, the non-negative orthant of R^s, and $q^0 \gg 0$. The appeal of allowing negative entries is clear, yet the boundedness from below is essential (as we shall see), therefore (iii) is not much of a generalization.

Assumptions (cont.)

(v) μ^* is a probability measure on T.
(vi) $F(x, \ell, t)$ is a Borel measurable function on $X \times L \times T$, and non-decreasing in the variable x.
(vii) $-\infty < v(\mu^*) < \infty$.

We introduce assumption (vii) to avoid trivial pathologies. The part $v(\mu^*) < \infty$ is implied, e.g., by the assumption [adopted in Groves and Hart (1982)] that $F(x, \ell, t) \leqq a + b|x|$ for some a and b.

For the next three assumptions we need the following notation: Let $G(t)$ be the set in R^{s+1}, defined by

$$G(t) = \text{closure } \{(x, \alpha) : x \in X, \alpha \leqq F(x, \ell, t) \text{ for some } \ell \in L\}.$$

We consider the collection of closed subsets of R^{s+1} as a metric space with the metric generated by closed convergence [see Hildenbrand (1974)], namely A_k converges to A if every cluster point of $\{a_k\}$, with $a_k \in A_k$, is in A, and every point $a \in A$ is a limit of a sequence a_k with $a_k \in A_k$.

Assumptions (cont.)

(viii) $G(t)$ has a bounded selection.
(ix) If t is an atom of μ^* then $G(t)$ is convex.
(x) The mapping $t \to G(t)$ is continuous.

The first of the previous three conditions amounts to normalization. The second is standard in the economic literature and it is not a serious restriction in our model. Indeed, when producers are sampled, and several of them represent the same atom, then in the limit we have the convexification effect; namely, in the limit the system behaves, as far as computing the value, as if it satisfies the assumption. We leave out these details and assume (ix). The difference may arise in the allocation process and we comment on this when discussing allocation schemes (Remark 6.5). Assumption (x) looks strong [in Groves and Hart (1982) the set T is merely a measure space], but actually it is not a restriction at all, and we adopt it here for convenience of presentation only. The reason is that the analysis can anyway be done on the space of agents characteristics; namely, what counts is not the name t, but rather the type as expressed in the production function $F(x,\ell,t)$, and any sample from T can be interpreted as a sample of production functions. Whatever natural topology is taken on the space of production functions, the mapping $t \to G(t)$ [or rather, $F \to G(F)$] will turn out to be continuous. Rather than introducing this machinery and prove the continuity at the level of the agents' characteristics, we prefer to assume it at the level of the producers' names.

Some notations: $p \cdot x$ denotes the scalar product of p and x; the norm $|x|$ is the Euclidean norm, i.e., $|x| = (x \cdot x)^{1/2}$; inequalities between vectors are interpreted coordinatewise. Domains of integration under an integral sign, and arguments of integrands, will be supressed when obvious.

3. The approach

Pursuing the motivation that if $N = \{t_1, \ldots, t_n\}$ is a good example of T then $v(\mu^*)$ is a good approximation of $v(N)$, Arrow and Radner (1979) came up with the following result. It has been established later by Groves and Hart (1982) under relaxed conditions.

Proposition 3.1. If $t_1, \ldots, t_n$ are drawn as iid (independently and identically distributed) with the common distribution being μ^*, then $v(N)$ converges almost surely to $v(\mu^*)$, as n tends to infinity.

Two conceptual drawbacks of the previous approach are apparent. The first is the probabilistic nature of the conclusion. The convergence holds only almost surely, and may fail at a given realization, or in a specific example. The second condition is the independence condition which seems very restrictive from the point of view of applications.

In this paper we suggest another approach. We consider the empirical distribution of the sample $N = \{t_1, \ldots, t_n\}$, and try to examine whether closeness of this empirical measure to μ^* guarantees closeness of $v(N)$ to $v(\mu^*)$. If, or rather when, the implication holds, we have convergence that is deterministic and does not depend on the process in which the sample is formed. We comment later on how to recover Proposition 3.1 from our results. This approach of, essentially, comparing the distributions of the agents' characteristics is of course not new in the economics literature; it is used extensively in general equilibrium theory of large economies, see Hildenbrand (1974).

We need the following notions: Given a sample $N=\{t_1,\dots,t_n\}$, we denote by μ_n the empirical measure generated by N, namely μ_n is a probability measure and $\mu_n(A)$ is equal to n^{-1} times the number of indices i for which $t_i \in A$. As a convergence notion among probability measures on T we take the weak convergence of measures, or equivalently, the Prohorov distance will serve as a metric, see Billingsley (1968, p. 11 and p. 238), or Hildenbrand (1974, p. 48).

With the previous conventions, our approach is to investigate the continuity of the values $v(N)$ with respect to the measures μ_n. To this end it may be natural not to restrict the argument to be an empirical measure, namely to examine the continuity of $v(\mu)$ on μ, and as μ converges to μ^*. Indeed, that is what we do in section 5. But before that, we show in the next section that what may appear intuitive is not correct; namely, without an additional condition, $v(N)$ may not converge to $v(\mu^*)$ as the empirical measures converge to the underlying distribution.

4. Counterexamples

We start with a simple example in which the empirical measures μ_n converge, yet the value of the limit economy is inferior to the limit values of the samples. We then provide modifications of the same example, this to eliminate possible false conjectures as to the source of the phenomenon.

Example 4.1. Only two types of producers participate, say $T=\{a,b\}$. There are no local decisions, and the production functions (from which the variable ℓ is absent) are $F(x,a)=2x$ and $F(x,b)=3x$. Here x is a scalar and $X=[0,\infty)$. We set $q^0=1$. Let the sample N have one producer of type b and $n-1$ producers of type a. The optimal solution of $P(N)$ is obviously to provide the producer of type b with all the resources, thus getting $v(N)=3$. The limit distribution μ^* is, however, concentrated on $\{a\}$, and thus $v(\mu^*)=2$. (Indeed, continuing the discussion in section 2, it would be foolish in our case to use an optimal solution of the limit economy in any of the samples.)

The lack of convergence in the previous counterexample should not be attributed to the fact that each of the samples has some positive weight far from the support of the limit distribution. Nor is the reason that the production function is linear. In the following counterexample the support of each of the samples is included in the support of the limit distribution, and the production function of each producer is bounded. Yet the value of the limit variational problem is not a good approximation of the limit of the values.

Example 4.2. Again x is scalar, $X=[0,\infty)$ and $q^0=1$. There are no local decisions. The space T is $(0,1]$ and $F(x,t)=(1-t/2)\min(x,\frac{1}{2}t^{-1/2})$. Let the limit distribution μ^* be the uniform distribution over T. Notice that there is a unique solution to $P(\mu^*)$, namely $x(t)=\frac{1}{2}t^{-1/2}$, and the optimal value is

$$v(\mu^*)=\int_0^1\left(1-\frac{t}{2}\right)\frac{1}{2}t^{-1/2}\,\mathrm{d}t=5/6.$$

Consider now a sequence of samples as follows. The sample $N=\{t_1,\ldots,t_n\}$ is defined by $t_1=(2n)^{-2}$ and $t_i=i/n$ for $i=2,\ldots,n$. The empirical measures converge clearly to μ^*. The optimal solution of $P(N)$ is $x(t_i)=0$ for $i>1$ and $x(t_1)=n$. The value is

$$v(N)=(1-n^{-2})n\cdot\frac{1}{n}$$

and as $n\to\infty$ the values $v(N)$ converge to 1, which is bigger than 5/6.

The previous counterexample can be further modified so that $F(x,t)$ is strictly concave in the x-variable. We leave out the details.

In both examples we have that $\lim v(N)$ is bigger than $v(\mu^*)$. We show in the next section that this reflects a general property. However, this property relies on the technical assumptions, in particular assumptions (viii) and (x). Without these the case $\lim v(N)<v(\mu^*)$ may occur.

What causes the discontinuity phenomenon in the examples is, roughly, that in the samples a small portion dominates in the efficient production, and it is small enough to be washed away in the limit. The condition in the next section eliminates such a possibility, and yields the convergence.

5. Convergence

We start with the definition of p-tightness, which is the condition that guarantees the convergence, and explain the measure theory and the economic interpretation of it. Then we state the convergence result, revisit the examples of last section in light of this result, and prove it. We conclude with some comments including recovery of the Arrow and Radner result. Although in the applications we are interested in a sequence of empirical measures, we give the definition and prove the result for general probability measures.

The following function will be used. For $p\in R^s$, a vector of prices, we set

$$S(t,p)=\sup\{F(x,\ell,t)-p\cdot x: x\in X, \ell\in L\}. \tag{5.1}$$

For any given measure on T, say v, the function $S(\cdot,p)$ is measurable with respect to the completion of the measure v. This is a standard application of the projection theorem, see Hildenbrand (1974, p. 44). In particular we can always integrate $S(t,p)$ against probability measures on T, with the integral, in view of (viii), being finite or $+\infty$.

Definition 5.1. Let $p\in R^s$ and let μ_j be a sequence of probability measures on T. We say that the sequence μ_j is p-tight if for every $\varepsilon>0$ there exists a compact subset T_ε of T such that $S(\cdot,p)$ is bounded on T_ε and

$$\int_{T\setminus T_\varepsilon} |S(t,p)|\,\mathrm{d}\mu_j<\varepsilon \quad \text{for all } j. \tag{5.2}$$

The p-tightness condition reflects an economic situation, and thus has an economic interpretation, as follows. The vectors p in the sequel are price vectors for the goods x. Given a vector p of prices, the amount $S(t,p)$ is the maximum (infinitesimal) profit that the producer t can make. If p-tightness holds, then although high profits by a small portion of the producers are possible, the effect of this on the whole economy is small; namely, in the aggregate, a small portion of the producers can with the given p, collect only a bounded portion of all profits.

The measure theory behind the definition can be understood best when the definition of tightness is recalled, e.g., Billingsley (1968, p. 37) or Hildenbrand (1974, p. 49). Indeed, p-tightness implies tightness of the measures v_j with $\mathrm{d}v_j(t)=|S(t,p)|\,\mathrm{d}\mu_j$. Alternatively, tightness is a special case of p-tightness where $S(t,p)=1$. Any convergent sequence of measures is tight, but may not be p-tight. One can give conditions on F that imply p-tightness for some or for all p. For instance, if $F(x,\ell,t)$ is bounded, or if $F(x,\ell,t)$ is $o\|x\|$ as $\|x\|\to\infty$, uniformly in (ℓ,t), then any converging sequence of measures is p-tight for all p in R^s_+.

For the statement of the main result we need to recall the notion of shadow prices for the underlying problem $P(\mu^*)$. If the latter has an optimal solution, say $(x^*(t),\ell^*(t))$, then (it is well-known), it generates prices, say p^*, such that

$$F(x^*(t),\ell^*(t),t)-p^*\cdot x^*(t)=S(t,p^*), \tag{5.3}$$

and (5.3) characterizes optimal solutions. But even when $P(\mu^*)$ does not have an optimal solution, the shadow prices p^* can be defined, with the property that there exists a maximizing sequence (namely a sequence of admissible policies with values converging to $v(\mu^*)$), say $(x_k(t),\ell_k(t))$, such that

$$F(x_k(t),\ell_k(t),t)-p^*\cdot x_k(t) \text{ converge } \mu^*-\text{a.e. to } S(t,p^*). \tag{5.4}$$

The geometrical interpretation of these shadow prices is that the vector $(-p^*,1)$ generates a supporting hyperplane to the closure of the set

$$\{(\int x(t)\,\mathrm{d}\mu^*, \int F(x(t),\ell(t),t)\,\mathrm{d}\mu^*):(x(t),\ell(t)) \text{ admissible}\}$$

at the point $(q^0, v(\mu^*))$; equivalently, if $v(\mu^*,q)$ denotes the supremum of the problem $P(\mu^*)$ with constraint condition q, instead of q^0, then

$$v(\mu^*)-p^*\cdot q^0=\max\{v(\mu^*,q)-p^*\cdot q: q\in X\}.$$

Theorem 5.2. Let μ_j be a sequence of probability measures on T, converging to the underlying probability measure μ^*. Let $v(\mu_j)$ be the value of the variational problem $P(\mu_j)$, then

$$\liminf v(\mu_j)\geqq v(\mu^*). \tag{5.5}$$

If in addition the sequence μ_j is p^*-tight, with p^* being the vector of shadow prices for $P(\mu^*)$, then

$$v(\mu_j) \text{ converge to } v(\mu^*). \tag{5.6}$$

Before turning to some lemmas and the proof of the theorem, we wish, in light of the convergence result, to examine the two examples of section 4.

Example 4.1, revisited. Firstly, we see a demonstration of conclusion (5.5), with strict inequality. The convergence (5.6) fails. Indeed, the shadow price for $P(\mu^*)$ is $p^*=2$. Any sample described in the example has positive weight on type b, and $S(b,2)=\infty$. Hence the 2-tightness fails.

Example 4.2, revisited. Again (5.5) holds with strict inequality. A possible shadow price for $P(\mu^*)$ is $p^*=\frac{1}{2}$, (actually any $0\leqq p^*<\frac{1}{2}$ will do). For $t_1=(2n)^{-2}$ we get $S(t_1,\frac{1}{2})=\frac{1}{2}n$, and when integrating $S(t,\frac{1}{2})$, as required in (5.2), on a set containing t_1, and with respect to μ_n, we get at least $\frac{1}{2}$. Since for large n the point t_1 is close to 0, and, eventually, outside of any compact set, we cannot achieve an arbitrary ε in (5.2) for a compact T_ε (although the integral is finite), and the sequence of empirical measures is not $\frac{1}{2}$-tight. [It is 2-tight, but 2 is not a shadow price for $P(\mu^*)$.] Notice that a slight change in the sampling would yield $\frac{1}{2}$-tightness. Indeed, if $t_1=1/n$, and as before $t_i=i/n$ for $i>1$, then $S(t_i,\frac{1}{2})=\frac{1}{4}(n/i)^{1/2}$, and the sequence μ_n is $\frac{1}{2}$-tight.

We prove the two parts of Theorem 5.2 separately, each part is preceded by some lemmas that are needed also in the sequel.

Lemma 5.3. Let ν be a measure on T, with compact support T_0. Let ν_j be a sequence of measures converging to ν. Let $f(t)$ be a bounded real-valued function, continuous at each $t_0 \in T_0$ (yet not necessarily continuous at $t \notin T_0$). Then $\int f\,d\nu_j$ converge to $\int f\,d\nu$.

Proof. If f were continuous on T the result would follow from weak convergence. So let g be bounded and continuous on T, such that g coincides with f on T_0. Such g exists by Tietze Theorem. Then $\int g\,d\nu_j$ converge to $\int g\,d\nu$. But in view of the Prohorov metric characterization of weak convergence [Billingsley (1968, p. 238)], and since both f and g are bounded and $f(t)-g(t)$ tends to 0 as $t \to T_0$, we can conclude that $\int (f-g)\,d\nu_j \to 0$ as $j \to \infty$. This completes the proof.

Lemma 5.4. Let T_0 be a compact subset of T such that $\mu^*(T_0) \geqq 1-\varepsilon$. Let μ_j converge to μ^*. Let $(x(t), \ell(t))$ be admissible choices such that: $x(t)$ and $F(x(t), \ell(t), t)$ are bounded, say by the bound b, and both continuous at $t \in T_0$. Then for large j the integrals $\int x(t)\,d\mu_j$ are within an $\varepsilon(3b+1)$ neighborhood of $\int x(t)\,d\mu^*$, and $\int F(x(t), \ell(t), t)\,d\mu_j$ are within $\varepsilon(3b+1)$ neighborhood of $\int F(x(t), \ell(t), t)\,d\mu^*$.

Proof. We can separate each μ_j to, say, $\nu_j + \sigma_j$, such that ν_j converge to μ_0^*, the restriction of μ^* to T_0, and $\sigma_j(T) < \varepsilon$. The separation can be done in view of (i) by using standard arguments. Then, by Lemma 5.3, $\int x\,d\nu_j$ and $\int F\,d\nu_j$ converge, respectively, to $\int x\,d\mu_0^*$ and $\int F\,d\mu_0^*$. The latter are within εb neighborhoods of $\int x\,d\mu^*$ and $\int F\,d\mu^*$, and since $\sigma_j(T) < \varepsilon$ the integrals with respect to ν_j are within $2\varepsilon b$ neighborhoods of $\int x\,d\mu_j$ and $\int F\,d\mu_j$, hence the estimates.

We say that the admissible policy $(x(t), \ell(t))$ is an ε-solution of $P(\mu^*)$ if $\int x(t)\,d\mu^* \leqq q^0 + (\varepsilon, \ldots, \varepsilon)$ and $v(\mu^*) - \int F(x(t), \ell(t), t)\,d\mu^* < \varepsilon$.

Lemma 5.5. Let $\varepsilon > 0$ be specified. There exists an ε-solution $(x(t), \ell(t))$ of $P(\mu^*)$ with the following properties: $x(t)$ and $F(x(t), \ell(t), t)$ are both bounded, say by b; there exists a compact set T_0 with $\mu^*(T_0) \geqq 1 - \varepsilon b^{-1}$ and such that $x(t)$ and $F(x(t), \ell(t), t)$ are continuous at each $t \in T_0$.

Proof. We start with an $\varepsilon/2$-solution of $P(\mu^*)$. It can be made bounded by replacing its tail with the bounded selection guaranteed by condition (viii). If the tail is chosen small enough the result is, say, an $\frac{3}{4}\varepsilon$-solution, bounded, say by $b-1$. We denote this solution by $(x_1(t), \ell_1(t))$. By Lusin theorem there exists a compact set T_0 with $\mu^*(T_0) \geqq 1 - \frac{1}{4}\varepsilon b^{-1}$ such that $x_1(t)$ and $F(x_1(t), \ell_1(t), t)$ are continuous on T_0. The continuity of $G(t)$, assumed in (x), enables to choose $x(t)$ and $\ell(t)$ on $T \backslash T_0$ such that $x(t)$ and $F(x(t), \ell(t), t)$ are bounded by b and such that $t \to t_0$ and $t_0 \in T_0$ then $\lim x(t) = x_1(t_0)$ and $\lim F(x(t), \ell(t), t) = F(x_1(t_0), \ell_1(t_0), t_0)$. This can be done by applying a standard measurable selection theorem at a set-valued function which pointwise is a subset of G and at each $t_0 \in T_0$ it is equal to $\{(x_1(t_0), F(x_1(t_0), \ell_1(t_0), t_0))\}$. Such a set-valued mapping can easily be constructed in view of the continuity of x_1 and $F(x_1, \ell_1, t)$ on T_0. The policy $(x_1(t), \ell_1(t))$ for $t \in T_0$ and $(x(t), \ell(t))$ for $t \in T \setminus T_0$ yields now the desired properties for x and F, namely boundedness by b and continuity at $t \in T_0$. The choice of T_0 implies that $(x(t), \ell(t))$ yields an output which differs in value from that of $(x_1(t), \ell_1(t))$ by only $\frac{1}{4}\varepsilon$, hence $(x(t), \ell(t))$ is an ε-solution, with the desired properties.

Proof of (5.5) in Theorem 5.2. The idea is to consider the ε-solution guaranteed in Lemma 5.5, and apply to it the conclusion of Lemma 5.4. This would show that $v(\mu^*) - \varepsilon$ can be obtained in $P(\mu_j)$ for j large, and since ε is arbitrary, (5.5) holds. There is, however, one flaw: The ε-solution may not satisfy the constraint q^0 for $P(\mu_j)$ or for $P(\mu^*)$. We therefore have to modify the argument, as follows.

The resources q^0 were fixed in the definition of $P(\mu)$. We can, however, consider the problem with various constraints q. We therefore write $P(\mu, q)$ for the variational problem with resources q, and write $v(\mu, q)$ for the supremum of this problem.

Consider now $P(\mu^*, q)$. We claim that $v(\mu^*, q)$ is a convex function of q. This follows in a standard way from Liapunov convexity theorem [Hildenbrand (1974, p. 45)] and condition (ix). Together with condition (iv) the convexity of $v(\mu^*, q)$ implies that $v(\mu^*, q)$ is continuous at q^0. Therefore, for a given $\varepsilon_0 > 0$ we can choose a $q^1 = q^0 - (\delta, \delta, \ldots, \delta)$ with $\delta > 0$, such that $v(\mu^*, q^0) - v(\mu^*, q^1) < \varepsilon_0$.

We now use the ε-solution guaranteed by Lemma 5.5 with $\varepsilon = \min(\delta, \varepsilon_0)$. The convergence that Lemma 5.4 implies, guarantees that for j large $\int x(t)\,d\mu_j$ is within a δ-neighborhood of q^1, in particular it satisfies the q^0 constraint. And that for j large

$$v(\mu^*, q^0) - \int F(x(t), \ell(t), t)\,d\mu_j < \varepsilon_0,$$

in particular $v(\mu^*, q^0) - v(\mu_j, q^0) < \varepsilon_0$. Since ε_0 is arbitrary, (5.5) holds.

We need two lemmas before proving the other conclusion.

Lemma 5.6. Let μ_j converge to μ^*, and let p^* be a vector of shadow prices for $P(\mu^*)$. If μ_j is p^*-tight then it is p-tight whenever $p \geqq p^*$ (namely $p_i \geqq p_i^*$ for all the coordinates $i = 1, \ldots, s$).

Proof. Since $S(t, p)$ is bounded from below [by (viii)] and since μ_j is tight, by the convergence, we can conclude the p-tightness from an economic reasoning. For higher cost, $p \geqq p^*$, the profit is less, $S(t, p) \leqq S(t, p^*)$.

Lemma 5.7. Let $p_0 \in R^s$ and $T_0 \subset T$ be such that $S(t,p)$ is bounded as p ranges in a neighborhood of p_0, and $t \in T_0$, then the restriction of $S(t,p_0)$ to T_0 is continuous.

Proof. Notice that $S(t,p) = \sup\{r - p \cdot x : (x,r) \in G(t)\}$. The continuity of $G(t)$ implies that $S(t,p_0)$ is lower semicontinuous, i.e., if $t_j \to t_0$ then $S(t_0,p_0) \leqq \liminf S(t_j,p_0)$. Continuity will therefore be established when we verify that $S(t_0,p_0) \geqq \limsup S(t_j,p_0)$. We shall assume the contrary and get a contradiction.

If $S(t_0,p_0) < \limsup S(t_j,p_0)$ then a subsequence of t_j, say t_j itself, exists and $(x_j,r_j) \in G(t_j)$, such that $S(t_0,p_0) < \lim(r_j - p_0 \cdot x_j)$. The sequence (x_j,r_j) cannot be then bounded, since a cluster point (x_0,r_0) then belongs to $G(t_0)$, and by continuity $S(t_0,p_0) \geqq r_0 - p_0 \cdot x_0$. By boundedness of $S(t,p_0)$ it follows then that $\{x_j\}$ is unbounded. But then we can consider $p_j = p_0 - \varepsilon x_j \cdot |x_j|^{-1}$, with $\varepsilon > 0$ fixed so that p_j is in the prescribed neighborhood of p_0, and clearly $S(t_j,p_j)$ is unbounded. This is the desired contradiction.

Proof of (5.6) in Theorem 5.2. If (5.6) fails, then, with (5.5), there exists a subsequence of μ_j, say μ_j itself, with $v(\mu^*) < \lim v(\mu_j)$. Denote the right-hand side by r_0. Let p^* be the shadow price of $P(\mu^*)$. Then

$$r_0 - p^* \cdot q^0 > v(\mu^*) - p^* \cdot q^0. \tag{5.7}$$

We plan to show that under the given p^*-tightness, (5.7) is impossible. If (5.7) holds then also

$$r_0 - p_0 \cdot q^0 > v(\mu^*) - p_0 \cdot q^0 \tag{5.8}$$

for $p_0 > p^*$ and close enough to p^*. [Here we use condition (iii). It is possible to construct a counterexample to (5.6) if (iii) fails.] It follows from the measurability conditions [see e.g., Hildenbrand (1974, p. 65) for a correspondence form of this] that

$$v(\mu^*) - p_0 \cdot q^0 = \int_T S(t,p_0)\,\mathrm{d}\mu^*.$$

We use now the p^*-tightness, and the p_0-tightness (Lemma 5.6), and construct for $\varepsilon = \frac{1}{2}(r_0 - v(\mu^*))$ a set $T_\varepsilon \subset T$ such that

$$\int_{T \setminus T_\varepsilon} S(t,p_0)\,\mathrm{d}\mu_j < \varepsilon \tag{5.9}$$

for all j and such that $S(t,p^*)$ is bounded on T_ε (see Definition 5.1). It is clear that $S(t,p)$ is bounded for $t \in T_\varepsilon$ also for p in a neighborhood of p_0 [by boundedness of $S(t,p^*)$]. By Lemma 5.7 $S(t,p_0)$ is continuous on T_ε. By the convergence of μ_j to μ^* and (5.9) it follows that

$$\limsup \int_T S(t,p_0)\,d\mu_j < r_0 - p\cdot q^0.$$

But the latter inequality contradicts the definition of

$$S(t,p_0) = \sup\{r - p_0\cdot x : r = F(x,\ell,t),\ x\in X, \ell\in L\}.$$

This implies that (5.8), hence (5.7), are impossible, and completes the proof.

Remark. The previous proof could be simpler had we known that $S(t,p^*)$ is continuous. Unfortunately it may not be continuous even if it is bounded. We need to pass to an interior vector of shadow prices to get the continuity.

A natural query now is for conditions guaranteeing the conditions of the convergence result, in particular the p^*-tightness. Boundedness of $F(x,\ell,t)$ clearly suffices for all p. If $F(x,\ell,t)\leqq\alpha+p\cdot x$ then p_1-tightness is guaranteed for all $p^1\geqq p$. But recall that p-tightness may depend strongly on the sequence μ_j (see the revisit at Example 4.2 in this section), or in other words, the p-tightness may depend on the way the sample is formed.

In the Arrow and Radner procedure the samples are formed as independent and identically distributed drawings from μ^*. It is well-known then, that the empirical measures converge almost surely to μ^*; in particular the empirical measures are a.s. tight. It is as easy to check that the empirical measures are also almost surely p^*-tight; in fact, the empirical measures are a.s. p-tight for every p such that $S(p,t)$ is μ^*-integrable. This almost sure p^*-tightness allows to derive the Arrow and Radner result from our theorem.

6. Allocation procedures

In this section we examine ways to employ solutions of $P(\mu^*)$ in getting approximate solutions to $P(N)$; this when the sample is good and the conditions of the convergence result hold. We work here under the following assumption, used in Groves and Hart (1982). (The assumption is eliminated in the next section.)

Assumption 6.1. The producers and the center know the production function $F(x,\ell,t)$ for all t and know the underlying probability measure μ^*.

The allocation pattern. With Assumption 6.1 the approximate solutions have the following pattern. A pair $(x_0(t),\ell_0(t))$ is announced, or computed by the agents. The pair does not depend on the specific sample N. The producer t_i makes the local decision $\ell_0(t_i)$ and demands $x_0(t_i)$ as an input.

The benefits from such a program are apparent. If $(x_0(t), \ell_0(t))$ is an optimal solution of $P(\mu^*)$, then each t_i, under Assumption 6.1, can compute $(x_0(t_i), \ell_0(t_i))$ for himself/herself, with no need of communication, i.e., with fully decentralized decisions. Even if $(x_0(t), \ell_0(t))$ is a modification of an optimal solution, it can be announced once by the center (if it is supported by prices then it is enough to announce the prices) and then the agents act independently and decisions are decentralized, thus saving the effort needed in establishing a particular optimal solution to $P(N)$.

In the sequel we do not analyze the advantages of the program, but rather concentrate on the possibility of finding such a pair $(x_0(t), \ell_0(t))$. To this end we use the following.

Definition 6.2. Let $N = \{t_1, \ldots, t_n\}$ be an element in a sequence of samples. The pair $(x_0(t), \ell_0(t))$ is an asymptotic solution of $P(N)$ if

$$\limsup_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} x_0(t_i) \leqq q^0 \tag{6.1}$$

and

$$\lim_{n \to \infty} \left(v(N) - \frac{1}{n} \sum_{i=1}^{n} F(x_0(t_i), \ell_0(t_i), t_i)\right) = 0. \tag{6.2}$$

The pair $(x_0(t), \ell_0(t))$ is an asymptotic ε-solution if (6.1) holds and

$$\limsup_{n \to \infty} \left(v(N) - \frac{1}{n} \sum_{i=1}^{n} F(x_0(t_i), \ell_0(t_i), t_i)\right) \leqq \varepsilon. \tag{6.3}$$

The program that we display is the one behind the procedures suggested by Groves and Hart (1982). Our definition of an asymptotic solution differs slightly. We require that the inequality in $P(N)$ holds asymptotically, as reflected in (6.1), while Groves and Hart demand that it holds for each $P(N)$. Both approaches have clear economic interpretations. To take care of the case when $\sum x_0(t_i)$ exceeds nq, Groves and Hart (1982) provide rationing rules. We note, without supplying the details, that these rationing rules can be used also under the deterministic framework of this paper.

Theorem 6.3. Suppose that $(x^*(t), \ell^*(t))$ is an optimal solution of $P(\mu^*)$, and *such that both $x^*(t)$ and $F(x^*(t),\ell^*(t),t)$ are bounded and continuous. Let N_j be a sequence of samplings, with empirical measures μ_j which converge to μ^*, and which are p^*-tight, p^* being shadow prices for $P(\mu^*)$. Then $(x^*(t),\ell^*(t))$ is an asymptotic solution of $P(N_j)$.*

Proof. Property (6.1) is taken care of by the convergence of μ_j to μ^*, which also implies that

$$\frac{1}{n}\sum_{i=1}^{n} F(x^*(t_i),\ell^*(t_i),t_i) \text{ converge to } \int F(x^*(t),\ell^*(t),t)\,\mathrm{d}\mu^*.$$

The latter is equal to $v(\mu^*)$. Property (6.2) follows now from Theorem 5.2 which implies that $v(N_j)$ converge to $v(\mu^*)$.

Unfortunately, the general conditions that we work under, do not imply existence of an optimal solution $(x^*(t),\ell^*(t))$, and even when it exists, it may not satisfy the conditions needed in the preceding theorem. In these cases $(x^*(t),\ell^*(t))$ cannot be used as an asymptotic solution, but the following result may be of use.

Theorem 6.4. Let N_j be a sequence of samplings, with empirical measures μ_j coverging to μ^*, and which are p^*-tight, where p^* being shadow prices of $P(\mu^*)$. Then for every $\varepsilon>0$ there is an asymptotic ε-solution $(x_0(t),\ell_0(t))$ of $P(N_j)$. This asymptotic ε-solution can be obtained by taking any $\varepsilon/3$-solution of $P(\mu^*,q_1)$ with $q_1<q^0$ and close enough to q^0, and modify it on a set T_1 of small μ^*-measure, such that at any $t\in T_0$ with $T_0=T\backslash T_1$, the functions $x_0(t)$ and $F(x_0(t),\ell_0(t),t)$ are continuous.

Proof. How to choose q_1 is described at the beginning of the proof of Theorem 5.2 (and then the strict inequality constraint holds for j large). How to construct then $(x_0(t),\ell_0(t))$ is described in Lemma 5.5, with the convergence of $(1/n)\sum F(x_0(t_i),\ell_0(t_i),t_i)$ to $v(\mu^*)-\varepsilon$ guaranteed in Lemma 5.4. The convergence result, Theorem 5.2, implies that $(x_0(t),\ell_0(t))$ is an asymptotic ε-solution.

Remark 6.5. The convergence result of the previous section, and the derived asymptotic solutions of the present section, make use of the convexity assumption (ix). We explained in the previous section why the convergence result holds when the μ_j are the empirical measures of a sequence $t_1,t_2,\ldots$ of agents that are drawn according to the underlying measure μ^*. In such a case it is also possible to modify the allocation pattern of the present section, and maintain the optimality without the convexity. This can be done as follows [a similar technique was used in Groves and Hart (1982)]. First we allow mixed strategies, or relaxed strategies in the languages of the calculus of variations. Namely, we allow an agent of type τ to consider pairs $(x_1,\ell_1),\ldots,(x_{s+2},\ell_{s+2})$ with, respectively, probabilities $\alpha_1,\ldots,\alpha_{s+2}$. The resources exploited by such a choice are interpreted as $\sum\alpha_i x_i$, and the output as $\sum\alpha_i F(x_i,\ell_i,\tau)$. The set $G(\tau)$ for the relaxed problem is convex; indeed, it is a subset of R^{s+1}, and by Caratheodory theorem convex combinations of $s+2$ elements $(x_i, F(x_i,\ell_i,\tau))$ generate the convex hull. The modified problem has an asymptotic ε-solution, and let $(x_0(t),\ell_0(t))$ be one. For the atom τ the choice $(x_0(\tau),\ell_0(\tau))$ may be the relaxed one, but then $(x_0(\tau),F(x_0(\tau),\ell_0(\tau),\tau))$ is generated as a convex combination of, say, $(x_{0,i},F(x_{0,i},\ell_{0,i},\tau))$, $i=1,\ldots,s+2$, with weights α_i, $i=1,\ldots,s+2$. When the agents arrive, a subsequence of them, say $t_{i_1},t_{i_2}\ldots$, is of type τ, with proportion $\mu^*(\tau)$ to the whole sequence. What ought to be arranged now in order to get an ε-asymptotic solution without convexity is that each t_{i_j} will choose one of the $(x_{0,i},\ell_{0,i})$ in a way that asymptotically each of $(x_{0,i},\ell_{0,i})$ is chosen with proportion α_i. This can be arranged in two ways. One mechanism is probabilistic: t_{i_j} chooses one of the $(x_{0,i},\ell_{0,i})$ with probabilities $\alpha_1,\ldots,\alpha_{s+2}$. The decisions are decentralized, and the goal is achieved with probability one. A deterministic procedure can be arranged, but needs some centralization: As the agents t_{i_j} arrive, they are assigned one of the $(x_{0,i},\ell_{0,i})$ by the center, such that the correct asymptotic proportions are obtained.

7. An adaptive procedure

The complete knowledge of Assumption 6.1 does not occur in many realistic situations. Often, the underlying probability μ^* exists but is not known beforehand, and information about it, and about the technologies $F(x,\ell,t)$, has to be collected during the allocation process. An adaptive scheme is then in place, and in this section we examine the following procedure.

An adaptive allocation pattern. As the sequence of producers $t_1,t_2,\ldots$, come with their demands, the center collects information about their production functions $F(x,\ell,t_i)$. At a given sequence of arrivals, say $i_1,\ldots,i_j,\ldots$, the center announces, respectively, prices, $p_{i_1},\ldots,p_{i_j},\ldots$. The producers t_k with $i_j<k\leq i_{j+1}$ choose then (x_k,ℓ_k) so to maximize $F(x_k,\ell_k,t_k)-p_{i_j}\cdot x_k$.

Definition 7.1. If the outcome of the allocation pattern is that

$$\limsup_{n\to\infty}\frac{1}{n}\sum_{k=1}^{n} x_k\leq q^0 \tag{7.1}$$

and

$$\lim_{n\to\infty}\left(v(N)-\frac{1}{n}\sum_{k=1}^{n}F(x_k,\ell_k,t_k)\right)=0, \tag{7.2}$$

then we say that $p_{i_1}, p_{i_2}, \ldots$ generate an asymptotic solution to $P(N)$ (compare with Definition 6.1).

In Remark 7.6 we comment on the efficiency of this procedure, and on alternative procedures. For the allocation pattern to work we need a technical assumption which is stronger than what we used before [yet not stronger than the assumption used by Arrow and Radner (1979) or the one used by Groves and Hart (1982) when dealing with a price mechanism].

Technical Assumptions 7.2. The function $g(x,t)=\sup\{F(x,\ell,t):\ell\in L\}$ is strictly concave in the variable x, and for each (x,t) there is an ℓ such that $g(x,t)=F(x,\ell,t)$. The problem $P(\mu^*)$ has an optimal solution.

The condition that $g(x,t)$ is attained can be relaxed, see Remark 7.6. Note that strict concavity in x of $F(x,\ell,t)$ is not sufficient for the technical assumptions. If, however, ℓ belongs to a vector space [as assumed in Arrow and Radner (1979)] and $F(x,\ell,t)$ is strictly concave in (x,ℓ), and $g(x,t)$ is attained, then the strict concavity of $g(x,t)$ follows. Conditions guaranteeing existence of solutions for $P(\mu^*)$ are available in the literature [e.g. Aumann and Perles (1965)].

Recall that p_n is a shadow price vector for the problem $P(N)$ if

$$v(N)-p_n\cdot q^0=\max\{v(N,q)-p_n\cdot q: q\in X\},$$

where $v(N,q)$ is the supremum of the problem $P(N)$ but with constraint q rather than q^0 (see the discussion preceding Theorem 5.2).

Theorem 7.3. Let $t_1, t_2, \ldots$ be a sequence of producers and suppose that the empirical measures μ_j, generated by $\{t_1,\ldots,t_j\}$ converge to μ^*. Suppose also that every shadow price vector p^* to $P(\mu^*)$ is in the interior of the set of price vectors p for which $\{\mu_j\}$ is p-tight. Let p_{i_j} be a shadow price vector for the problem $P(\mu_{i_j})$. Then $p_{i_1}, p_{i_2}, \ldots$ generate an asymptotic solution to $P(N)$.

Proof. We first argue that $\{p_{i_j}\}$ is a bounded sequence, and every cluster point, say p_0, is a shadow price for $P(\mu^*)$. This holds under more general conditions; in our framework it follows directly from the observation that $v(N,q)$ is a concave function and converges to $v(\mu^*,q)$ on a set of q in a neighborhood of q^0. The concavity of $v(N,q)$ follows from the concavity of $g(x,t)$, and the convergence follows from Theorem 5.2 and the condition that μ_j are p-tight for a neighborhood of the shadow prices for $P(\mu^*,q^0)$.

The strict concavity of g implies that the demand coordinate $x^*(t)$ of all possible solutions $(x^*(t), \ell^*(t))$ of $P(\mu^*)$, is unique; indeed it is determined by the condition

$$\text{maximize } g(x,t) - p^* \cdot x$$

with p^* any shadow price vector of $P(\mu^*)$. Furthermore, $x^*(t)$ is continuous, and $F(x^*(t), \ell^*(t), t)$ is continuous. This follows from the strict concavity of $g(x,t)$ and the continuity assumption (x); indeed, the pairs $(x, g(x,t))$ form the boundary of $G(t)$.

Let p_j be a shadow price vector to $P(\mu_j)$, and denote by $y_j(t)$ the solution of

$$\text{maximize } g(y,t) - p_j \cdot y.$$

Then, we claim, $y_j(t)$ converge to $x^*(t)$ and $g(y_j(t),t)$ converge to $F(x^*(t), \ell^*(t), t)$; furthermore, the convergence is uniform on compact sets. The convergence follows from the first observation of the proof, namely that p_j converge to the price vectors that characterize $x^*(t)$, and from the strict concavity of $g(x,t)$. The uniformity of the convergence on a compact T_0 follows from the continuity assumption (x), since the latter implies that $y_j(t_j)$ converge to $x^*(t)$ if $t_j \to t$.

Suppose we manage to find, for a given $\varepsilon > 0$, a compact T_0 such that on the complement of T_0 all the integrals of $y_k(t)$, $g(y_k(t),t)$, $x^*(t)$, $F(x^*(t), \ell^*(t), t)$ with respect to all μ_j and μ^*, are less than ε. In the sequel we say 'ε-converge to', and mean 'converge to an ε-neighborhood of'.

A consequence of the preceding observations is that for the assumed T_0

$$\int_{T_0} y_j(t)\,\mathrm{d}\mu_j \quad \varepsilon\text{-converge to} \quad \int_{T_0} x^*(t)\,\mathrm{d}\mu^* \tag{7.3}$$

and

$$\int_{T_0} g(y_j(t),t)\,\mathrm{d}\mu_j \quad \varepsilon\text{-converge to} \quad \int_{T_0} F(x^*(t), \ell^*(t), t)\,\mathrm{d}\mu^*. \tag{7.4}$$

This follows from the weak convergence of μ_j to μ^*, and the established continuity of $x^*(t), F(x^*(t), \ell^*(t), t)$ and the uniform convergence of $y_j(t), g(y_j(t),t)$ and the properties of T_0.

Notice now that as p_{i_j} are announced, the choices (x_k, ℓ_k) of the producers t_k with indices $i_j < k \leqq i_{j+1}$, are $x_k(t) = y_{i_j}(t)$. Therefore (7.3) and (7.4) with the uniform convergence imply

$$\frac{1}{n} \sum_{\substack{1 \leqq k \leqq n \\ t_k \in T_0}} x_k \quad \varepsilon\text{-converge to} \quad \int_{T_0} x^*(t)\,\mathrm{d}\mu^* \tag{7.5}$$

and

$$\frac{1}{n}\sum_{\substack{1\leq k\leq n\\ t_k\in T_0}} F(x_k,\ell_k,t_k) \quad \varepsilon\text{-converge to} \quad \int_{T_0} F(x^*(t),\ell^*(t),t)\,\mathrm{d}\mu^*. \tag{7.6}$$

Since integration on the complement of T_0 would change the integrals by at most ε, and if indeed ε is arbitrarily small (namely T_0 can be such constructed) we may conclude, that

$$\frac{1}{n}\sum_{k=1}^{n} x_k \quad \text{converge to} \quad \int_T x^*(t)\,\mathrm{d}\mu^* \tag{7.7}$$

and

$$\frac{1}{n}\sum_{k=1}^{n} F(x_k,\ell_k,t_k) \quad \text{converge to} \int_T F(x^*(t),\ell^*(t),t)\,\mathrm{d}\mu^*. \tag{7.8}$$

However, $(x^*(t),\ell^*(t))$ is an optimal solution of $P(\mu^*)$, therefore (7.7) implies (7.1) and (7.2) is implied by (7.8) together with the convergence of $v(N)$ to $v(\mu^*)$, which is the content of Theorem 5.2.

The proof will be complete once we establish the existence of the promised set T_0. This existence is implied by the p-tightness condition, as follows.

Let $r_1,\dots,r_m$ be price vectors such that $\{\mu_j\}$ is r_i-tight for all $i=1,\dots,m$ and the shadow prices of $P(\mu^*)$ are contained in the interior of the convex hull of $\{r_1,\dots,r_m\}$, which we denote by R. The concavity of $g(x,t)$ and condition (viii) imply that the sets T_ε guaranteed by the p-tightness (see Definition 5.1) can be chosen the same for all $p\in R$. Since μ_j converges, in the weak convergence of measures, if follows that T_ε can be chosen so that $\int_{T_\varepsilon} c\,\mathrm{d}\mu_j$ is uniformly small, with c a constant. We choose the constant c such that $S(t,p)+c$ is positive for $p\in R$[such a c exists by (viii)]. The p-tightness can then be expressed as follows. For every $\delta>0$ there exists a compact T_δ such that

$$\int_{T\setminus T_\delta} |S(t,p)|\,\mathrm{d}\mu_j<\delta/2 \quad \text{and} \quad \int_{T\setminus T_\delta} c\,\mathrm{d}\mu_j<\delta/2 \tag{7.9}$$

for all j and $p\in R$.

Let k be so large that an $\eta>0$ exists such that p_k+p is in R if $|p|\leqq\eta$ [here p_k is the shadow price vector of $P(\mu_k)$]. Recall that

$$S(t,p_k)=g(y_k(t),t)-p_k\cdot y_k(t) \tag{7.10}$$

and similarly with $x^*(t),p^*$, replacing $y_k(t)$ and p_k. We claim that

$$\left| \int_{T \setminus T_\delta} y_k(t)\, \mathrm{d}\mu_j \right| < 2\eta^{-1}\delta \tag{7.11}$$

for all j. Otherwise a vector p of norm η can be found such that

$$\int_{T \setminus T_\delta} p \cdot y_k(t)\, \mathrm{d}\mu_j \geqq 2\delta,$$

hence, by (7.10) and (7.9),

$$\int_{T \setminus T_\delta} (g(y_k(t), t) - (p_k - p)\, y_k(t))\, \mathrm{d}\mu_j \geqq \delta, \tag{7.12}$$

but the integrand in the former expression is smaller than $S(t, p_k - p)$, a contradiction to (7.9).

Once (7.11) holds, (7.10) and (7.9) imply that

$$\left| \int_{T \setminus T_\delta} g(y_k(t), t)\, \mathrm{d}\mu_j \right| \leqq (2\eta^{-1} + 1)\delta \tag{7.13}$$

and, similarly, (7.11) and (7.13) hold with $x^*(t)$ replacing $y_k(t)$.

Given $\varepsilon > 0$, the δ can be arranged such that $(2\eta^{-1} + 1)\delta \leqq \varepsilon$, then T_δ can serve as the desired T_0. This completes the proof of the theorem.

We produce now two examples, the first shows that the strict concavity of $g(x, t)$ cannot be dropped, the second shows that p-tightness for an open set around p^* cannot be dropped.

Example 7.4. Let $T = [0, 1]$, $X = [0, \infty)$ and $q^0 = 1$. Let μ^* be supported on the Cantor set C in T. For $t \in C$ let $F(x, t)$ (the variable ℓ is absent) be given by $F(x, t) = x^{1/2}$ if $x \leqq 1$, $F(x, t) = \min(\frac{1}{2} + \frac{1}{2}x, 2)$ if $x \geqq 1$. For t not in C let $F(x, t) = x^\alpha$ if $x \leqq 1$ and $F(x, t) = \min(1 - \alpha + \alpha x, 2)$ for $x \geqq 1$, with $\alpha = \alpha(t)$ such that $\alpha(t) < \frac{1}{2}$ and $\alpha(t)$ is continuous. The unique optimal solution to $P(\mu^*)$ is then $x(t) = 1$ for $t \in C$, and the shadow price p^* is 2. If, however, μ_j is supported outside the Cantor set then $p_j > 2$, p_j being the shadow price for $P(\mu_j)$. The sequence t_j can be arranged so that μ_j satisfies all the conditions of the theorem, but supported out of C; then x_k converge to 2, which would violate the constraint (7.1).

Example 7.5. Let $T = [0, 1)$, $X = [0, \infty)$ and $q^0 = 1$. Let μ^* be supported at $\{0\}$. For t in a neighborhood of $t = 0$ let $F(x, t) = x^\alpha$, with $\alpha = \alpha(t)$ strictly decreasing and $\alpha(0) = \frac{1}{2}$. For t in a neighborhood of 1, let $F(x, t) = 1 - \alpha + \alpha x - \mathrm{e}^{-x}$ for $x \leqq b(t)$, and $b(t)$ will be determined later on. The sequence t_i is such that $t_i \to 0$ except for a subsequence i_j such that μ_j still converge to μ^*. The subsequence t_{i_j} will converge very fast to 1 so that the following phenomenon happens. For $P(\mu^*)$ the unique shadow price is $p^*=2$. Yet for each μ_j the shadow price p_j is greater than 2. The element t_{i_j} is chosen such that x_{i_j} is very large, say $x_{i_j}=2i_j$. This can be arranged with $b(t)$ finite, so that x_{i_j} is well defined. Again the constraint (7.1) is not valid. Note that $\{\mu_j\}$ is 2-tight, but it is not $(2+\varepsilon)$-tight for all $\varepsilon>0$.

Remark 7.6. The condition that $g(x,t)$ is attained can be relaxed, and with it the condition that $P(\mu^*)$ has a solution. It is enough to assume that the problem of maximizing $\int g(x(t),t)\,\mathrm{d}\mu^*$ subject to $\int x(t)\,\mathrm{d}\mu^*=q^0$ has a solution. But then the allocation process has to be slightly modified, with more central decisions, as follows. At each i_j the center, in addition to p_{i_j}, announces an ε_j such that $\varepsilon_j\to 0$. The producers that are not able to maximize $g(x,t_k)-p_{i_j}\cdot x$, have to get an ε_j approximation of this quantity. It is clear that these small errors will be washed out in the limit.

Another modification of the process is that the center, instead of announcing prices, announces choices $(x_j(t),\ell_j(t))$ which the producers must follow. This will probably allow more eased technical conditions. We did not examine this possibility, neither have we considered a rationing option to overcome the difficulty exhibited in Examples 7.4 and 7.5

It may seem that in the process of producing the sequence of prices p_{i_j} the center has to solve an infinite number of optimization problems $P(N)$, an effort which we wanted to spare. Indeed, some central effort ought to be done by the center, but not that much. To obtain p_n it is enough to compute

$$G_n=n^{-1}(G(t_1)+\cdots+G(t_n))$$

(here $A+B=\{a+b\colon a\in A, b\in B\}$), and find a supporting vector of the form $(-p_n,1)$ to a point (q^0,r). Since G_n can be recursively computed by $G_{n+1}=(1/(n+1))G(t_{n+1})+(n/(n+1))G_n$, the memory needed by the center is limited to two sets only. As Arrow and Radner (1979) noted, in practice, say when production functions are Cobb–Douglas, only a limited number of parameters is required to describe these sets.

References

Arrow, K.J. and R. Radner, 1979, Allocation of resources in large teams, Econometrica 47, 361–385.

Aumann, R.J. and M. Perles, 1965, A variational problem arising in economics, Journal of Mathematical Analysis and Applications 11, 488–503.

Billingsley, P., 1968, Convergence of probability measures (Wiley, New York).

Groves, T. and S. Hart, 1982, Efficiency of resource allocation by uniformed demand, Econometrica 50, 1453–1482.

Hildenbrand, W., 1974, Core and equilibria of a large economy (Princeton University Press, Princeton, NJ).