Microeconomic Foundations I
Choice and Competitive Markets

Student's Guide

Chapter 16: General Equilibrium, Time, and Uncertainty

Summary of the Chapter

In this chapter, the model and results of general equilibrium, the stuff of Chapters 14 and 15, are adapted to account for time and uncertainty. The chapter begins with a discussion of how the all-markets-at-once framework of Chapters 14 and 15 can do this, if one considers as different commodities the same physical good or service consumed (or used as input or produced) at different times or in different contingencies. But then we look at economies with a sequence of markets, some meeting and clearing before others begin to function.

Models of this sort of dynamic economy can be divided roughly into two groups. In some, opportunties to move wealth across time and states of nature—which is accomplished by trading in various securities—are *complete*: Subject to one unifying budget constraint, any shifting of wealth is feasible. The equilibria of these complete-market economies provide outcomes identical to those provided by Walrasian equilibria of all-markets-at-once economies. The chapter provides a very complete theory of these.

But in other cases, opportunities to shift wealth across time and states of nature are insufficient to achieve every desired shift (again, subject to a budget constraint). These are *incomplete-market economies/equilibria*, and they admit many "nasty" phenomena. The chapter doesn't offer anything like a complete treatment of incomplete markets; in fact, all it really does is provide a couple of examples to show the sorts of nasties that occur.

This is, by page count, the longest chapter in the book, and it certainly has one of the more complex "plot lines." It also comes with one of the most complex (and sometimes confusing) systems of notation. *But, conceptually, it is not at all difficult.* If you

Copyright © David M. Kreps, 2011. Permission is freely granted for individuals to print single copies of this document for their personal use. Instructors in courses using *Microeconomic Foundations I: Choice and Competitive Markets* may print multiple copies for distribution to students and teaching assistants, or to put on reserve for the use of students, including copies of the solution to individual problems, if they include a full copyright notice. For any other use, written permission must be obtained from David M. Kreps

read it carefully and slowly, keeping track as you go with the elaborate system of notation, you should find it reasonably straightforward. That said, you will probably need to take some breaks: Good places to take a break are between Sections 16.3 and 16.4, and then between 16.5 and 16.6.

The plot line is:

- Section 16.1 provides the setting, a time/uncertainty/information structure.
- In Section 16.2, the notion of a commodity is enriched to include the *contingency* (a time and state-of-information) in which the commodity is consumed (or otherwise used or produced by a firm), and the general equilibrium constructions of Chapters 14 and 15 are adapted to this enriched notion of a commodity. In this adaptation, we imagine that markets in all (contingent) commodities are conducted at a single point in time, before any production or consumption takes place.
- Section 16.3 gives a first cut at a more "dynamic" view of markets. Here, there are spot markets in basic commodities in each contingency and, in addition, a set of markets at the outset in so-called *Arrow-Debreu contingent claims*.
- Section 16.4 is the heart of the chapter. Here, more complex "securities" than simple Arrow-Debreu contingent claims are formulated and included in an equilibrium framework. Characterizations of equilibria are provided, and results are given concerning when the equilibrium consumption allocations in a dynamic economy are identical with the Walrasian-equilibrium allocations of the economy of Section 16.2
- In Section 16.4, securities pay *financial dividends*, dividends denominated in the numeraire (or in units of account). This gives rise to a massive indeterminancy in securities prices. So in Section 16.5, we "redo" the analysis of Section 16.4 but under the assumption that securities pay their dividends in contingent commodities. (We limit analysis to the case where all dividends are paid in one basic commodity, in each contingency.)
- The focus in Sections 16.3 through 16.5 is on "complete-markets" equilibria, although a number of results are proved that are more general. In Section 16.6, we briefly discuss what can happen when markets are "incomplete," illustrating with a very simple example.
- Firms are omitted from all the analysis of dynamic economies, from Section 16.3 through 16.6. Section 16.7 concludes the chapter by discussing how firms can be brought back into the story, including in that story the idea that amongst the securities being traded are shares of equity in the firms. When markets are competitive and complete, no difficulties arise. But when they are not complete, extreme pathologies can occur, which we illustrate with a final example. The nature of these pathologies suggests that the tools developed in this volume are inadequate to model important economic phenomenon, which is where this volume ends.

Solutions to Starred Problems

■ 16.2. (a) Assumption 16.1 implies 16.1′. If a consumer is globally insatiable in f_t wheat, then she is clearly globally insatiable in f_t consumption.

(b) Here is the modified Proposition 16.4.

Proposition 16.4′. *Assumption 16.1′ holds. Suppose $(p, \boldsymbol{x})$ is a Walrasian equilibrium for the all-at-once market structure. Then $(r, q, \boldsymbol{x}, \boldsymbol{y})$ where*

$$r_{if_t} = p_{if_t}, \quad q_{f_t} = 1, \quad \textit{and} \quad \boldsymbol{y}^h_{f_t} = p_{f_t} \cdot (\boldsymbol{x}^h_{f_t} - e^h_{f_t})$$

is an EPPPE for the economy with contingent financial markets. And if $(r, q, \boldsymbol{x}, \boldsymbol{y})$ is an EPPPE for the economy with contingent financial markets, then $(p, \boldsymbol{x})$ where

$$p_{f_0} = r_{f_0} \quad \textit{and} \quad p_{if_t} = q_{f_t} r_{if_t}, \textit{for } f_t \neq f_0$$

is a Walrasian equilibrium for the all-at-once market structure.

(c) At the risk of overdoing things, I reproduce the proof of Proposition 16.4, modified appropriately, below. Commentary and changes are provided in a sans serif font. But before launching into the proof, I repeat here the new budget constraint (for financial assets), giving them an "equation" number. The budget constraint for f_0 doesn't change; what are different (and simpler) are the constraints for $f_t \neq f_0$:

$$r_{f_0} \cdot x^h_{f_0} + q \cdot y^h \le r_{f_0} \cdot e^h_{f_0}, \quad \text{and} \quad r_{f_t} \cdot x_{f_t} \le r_{f_t} \cdot e^h_{f_t} + y_{f_t} \text{ for } f_t \neq f_0. \tag{16.1′}$$

Proof. [The proof in the text begins with some fluff, which I omit.] Suppose that $(p, \boldsymbol{x})$ is a Walrasian equilibrium. Define r, q, and $\boldsymbol{y}$ as shown [which is a bit changed]. To show that $(r, q, \boldsymbol{x}, \boldsymbol{y})$ is an EPPPE, we must verify that each consumer satisfies her budget constraints with $\boldsymbol{x}^h$ and $\boldsymbol{y}^h$ and is maximizing utility subject to those budget constraints, and that markets clear.

Concerning the budget constraints, since r is defined to be p, the definition of $\boldsymbol{y}$ can be rewritten as $\boldsymbol{y}^h_{f_t} = r_{f_t} \cdot (\boldsymbol{x}^h_{f_t} - e^h_{f_t})$ [simpler than before]. Therefore, by definition, the second "half" of (16.1′) [change here] holds with equality: $\boldsymbol{y}^h$ is defined to make this so. As for the first half, begin with the budget constraint in the Walrasian equilibrium, which holds with equality because all consumers are locally insatiable:

$$p \cdot \boldsymbol{x}^h = p \cdot e^h.$$

Break each sum into terms for contingency f_0 and for all others:

$$p_{f_0} \cdot \boldsymbol{x}^h_{f_0} + \sum_{f_t \neq f_0} p_{f_t} \cdot \boldsymbol{x}^h_{f_t} \quad = \quad p_{f_0} \cdot e^h_{f_0} + \sum_{f_t \neq f_0} p_{f_t} \cdot e^h_{f_t}.$$

Rearrange terms and substitute r for p (since they are identical), to get

$$r_{f_0} \cdot \boldsymbol{x}^h_{f_0} + \sum_{f_t \neq f_0} r_{f_t} \cdot (\boldsymbol{x}^h_{f_t} - e^h_{f_t}) \quad = \quad r_{f_0} \cdot e^h_{f_0}.$$

Substitute $\boldsymbol{y}^h_{f_t}$ for $r_{f_t} \cdot (\boldsymbol{x}^h_{f_t} - e^h_{f_t})$ in the summation and, recalling that $q \equiv 1$, we have the first half of (16.1').

Next is to show that the consumer is maximizing. Let $\hat{x}$ be any consumption bundle for h that, together with a corresponding $\hat{y}$, satisfies the budget constraints (16.1′). Sum up over all f_t the budget constraints for $\hat{x}$ and $\hat{y}$, and you get

$$r_{f_0} \cdot \hat{x}_{f_0} + q \cdot \hat{y} + \sum_{f_t \neq f_0} r_{f_t} \cdot \hat{x}_{f_t} \leq r_{f_0} \cdot e^h_{f_0} + \sum_{f_t \neq f_0} (r_{f_t} \cdot e^h_{f_t} + \hat{y}_{f_t}).$$

Note that $q \cdot \hat{y}$ on the left-hand side is the same as the summation of $\hat{y}$ terms on the right-hand side, since $q \equiv 1$, so these terms can be dropped. Then if you substitute ps for rs (they are identical), you get $p \cdot \hat{x} \leq p \cdot e^h$. Hence, $\hat{x}$ is a budget-feasible bundle for h in the economy with an all-at-once market structure and prices p. But since $\boldsymbol{x}^h$ is a Walrasian-equilibrium allocation at those prices, $u^h(\boldsymbol{x}^h) \geq u^h(\hat{x})$. That proves the maximization part of the definition of an EPPPE.

And to show that markets clear: The first half of (16.2) is the market-clearing condition for a Walrasian equilibrium (albeit expressed a bit less succinctly than usual), so it holds immediately. And, we know that $p \geq 0$ and $\sum_h \boldsymbol{x}^h \leq \sum_h e^h$, so clearly $p \cdot \sum_h (\boldsymbol{x}^h - e^h) \leq 0$. Moreover, this dot product consists only of nonpositive terms. And, because consumers are all locally insatiable, each satisfies Walras' Law with equality in any equilibrium, so this dot product is zero, and hence each term in it is zero. *That is, for each f_t,*

$$0 = p_{f_t} \cdot \sum_h \left(\boldsymbol{x}^h_{f_t} - e^h_{f_t}\right) = \sum_h p_{f_t} \cdot \left(\boldsymbol{x}^h_{f_t} - e^h_{f_t}\right) = \sum_h \boldsymbol{y}^h_{f_t}.$$

That is market clearing in the futures market; we know that $(r, q, \boldsymbol{x}, \boldsymbol{y})$ is an EPPPE.

For the other half, suppose that $(r, q, \boldsymbol{x}, \boldsymbol{y})$ is an EPPPE. To show that $(p, \boldsymbol{x})$ is a Walrasian equilibrium (for p defined as indicated from r and q), we have to show that each consumer is satisfying her (all-at-once) budget constraint, each is maximizing subject to that budget constraint, and that markets clear.

Market clearing is immediate from the first half of (16.2).

For the rest, we first argue that in any EPPPE, each q_{f_t} *[text eliminated here]* will be strictly positive. *Suppose to the contrary that, in an EPPPE, some $q_{f_t} \leq 0$. Temporarily denote the consumer whose preferences are globally insatiable in f_t by $\hat{h}$. We know then that* *there is some x' which is identical to $\boldsymbol{x}^{\hat{h}}$ (her equilibrium allocation) off of f_t and which satisfies $u^{\hat{h}}(x') > u^{\hat{h}}(\boldsymbol{x}^{\hat{h}})$. And she can afford x'; her budget constraints off of f_t are all satisfied with her original asset positions, since x' agrees with $\boldsymbol{x}^{\hat{h}}$ off of f_t. And to meet her f_t budget constraint, she simply buys enough financial claims for contingency-f_t numeraire as is required. Since $q_{f_t} \leq 0$, this at least maintains her f_0 budget constraint (if $q_{f_t} = 0$) and may even loosen it (if $q_{f_t} < 0$). In either case, x' becomes affordable, contradicting the optimality (for her) of $\boldsymbol{x}^{\hat{h}}$. This contradiction implies that q_{f_t} must be strictly positive for all f_t.*

Now take the EPPPE budget constraint for consumer h, given by (16.1′) *[change here, and...] Multiply the constraint for $f_t \neq f_0$ by q_{f_t}, and add all these constraints together (including the constraint for f_0), to get*

$$r_{f_0} \cdot \boldsymbol{x}^h_{f_0} + q \cdot \boldsymbol{y}^h + \sum_{f_t \neq f_0} q_{f_t} r_{f_t} \cdot \boldsymbol{x}^h_{f_t} \leq r_{f_0} \cdot e^h_{f_0} + \sum_{f_t \neq f_0} q_{f_t} \left(r_{f_t} \cdot e^h_{f_t} + \boldsymbol{y}^h_{f_t} \right).$$

Replace the (normalized) rs with ps, and cancel the common term $q \cdot \boldsymbol{y}^h$, and you get $p \cdot \boldsymbol{x}^h \leq p \cdot e^h$. So $\boldsymbol{x}^h$ is budget feasible for h at the prices p.

Finally, suppose that $\hat{x}$ is feasible for h at the prices p. Define $\hat{y} \in R^{N-1}$ by $\hat{y}_{f_t} = r_{f_t} \cdot (\hat{x}_{f_t} - e^h_{f_t})$ *[Change!]*. I assert that the pair $(\hat{x}, \hat{y})$ satisfies (16.1′), so that $\hat{x}$ is a feasible consumption bundle for h in the economy with futures markets and prices (r, q). *[Here and everywhere else that (16.1') appears, this is a change from before. I won't mention this again.]* Since $\boldsymbol{x}^h$ is an equilibrium bundle (in the EPPPE), once this is shown, we know that $u^h(\boldsymbol{x}^h) \geq u^h(\hat{x})$, proving that $\boldsymbol{x}^h$ solves h's utility maxmization problem in the all-markets-at-once economy, concluding the proof that $(p, \boldsymbol{x})$ is a Walrasian equilibrium for that economy.

Note that $\hat{y}$ is defined so that, for all f_t other than f_0, the second part of (16.1′) holds with equality. We only need to verify the first part of (16.1′). Since $\hat{x}$ is budget feasible in the all-markets-at-once economy at prices p, we know that $p \cdot \hat{x} \leq p \cdot e^h$. Write this as

$$p_{f_0} \cdot \hat{x}_{f_0} + \sum_{f_t \neq f_0} p_{f_t} \cdot \hat{x}_{f_t} \leq p_{f_0} \cdot e^h_{f_0} + \sum_{f_t \neq f_0} p_{f_t} \cdot e^h_{f_t},$$

hence

$$p_{f_0} \cdot \hat{x}_{f_0} + \sum_{f_t \neq f_0} p_{f_t} \cdot \left(\hat{x}_{f_t} - e^h_{f_t} \right) \leq p_{f_0} \cdot e^h_{f_0}.$$

In this inequality, replace each p with its definition in terms of q and r, and use the definition of $\hat{y}$, and you get *[the following equation is changed]*

$$r_{f_0} \cdot \hat{x}_{f_0} + \sum_{f_t \neq f_0} q_{f_t} r_{f_t} \cdot \left(\hat{x}_{f_t} - e^h_{f_t} \right) = r_{f_0} \cdot \hat{x}_{f_0} + \sum_{f_t \neq f_0} q_{f_t} \hat{y}_{f_t} = r_{f_0} \cdot \hat{x}_{f_0} + q \cdot \hat{y} \leq r_{f_0} \cdot e^h_{f_0}.$$

That is the first half of (16.1′), completing the proof. ■

■ 16.5. Recall the notation $\mathbf{S}(f_t)$ for the set of all successors to f_t.

(a) First, I'll give an arbitrage-based argument for the "qualitative" result that the price of a security is strictly positive at f_t if and only if it pays a strictly positive dividend in at least one successor to f_t:

Suppose that, for viable $\mathcal{S}$ and q, there is some $s \in \mathcal{S}$ and some $f_t \in \mathcal{F}(s)$ such that $q_{sf_t} > 0$, but $d_s(f'_{t'}) = 0$ for all $f'_{t'} \in \mathbf{S}(f_t)$. Then the plan of selling one unit of s in contingency f_t and holding to the end generates a wealth-transfer vector that is q_{sf_t} in contingency f_t and zero in every other contingency. This is an arbitrage opportunity, which is not allowed since $\mathcal{S}$ and q are viable. Hence, if $q_{sf_t} > 0$ for any $f_t \in \mathcal{F}(s)$, there must be some successor contingency to f_t in which s pays a strictly positive dividend.

On the other hand, suppose that we have $s \in \mathcal{S}$ and $f_t \in \mathcal{F}(s)$ such that $q_{sf_t} \le 0$ and, in some successor contingency to f_t, s pays a strictly positive dividend. Then the plan of buying a unit of s in f_t and holding to the end generates $-q_{sf_t}$ in f_t, nonnegative dividends in all successors to f_t, and a strictly positive dividend in at least one successor to f_t. That's an arbitrage opportunity. But this can't be if $\mathcal{S}$ and q are viable: If they are viable, and if, for $f_t \in \mathcal{F}(s)$, there is some successor contingency in which s pays a strictly positive dividend, then q_{sf_t} must be strictly positive.

Now I'll prove the more exact result, that if $\mathcal{S}$ and q are viable, then for all s and for all $f^0_\tau \in \mathcal{F}(s)$,

$$q_{sf^0_\tau} = \frac{1}{\pi_{f^0_\tau}} \sum_{f_t \in \mathbf{S}(f^0_\tau)} \pi_{f_t} d_s(f_t), \tag{16.4}$$

for all $\pi \in (\mathcal{S}, q)$:

Suppose $\mathcal{S}$ and q are viable. Fix $s \in \mathcal{S}$ and some $f^0_\tau \in \mathcal{F}(s)$. Consider the following trading plan: In contingency f^0_τ, buy one unit of s; then hold it to the end. This trading plan generates the wealth-transfer vector

$$\xi(f_t) = \begin{cases} -q_{sf^0_\tau}, & \text{if } f_t = f^0_\tau, \\ d_s(f_t), & \text{if } f_t \in \mathbf{S}(f^0_\tau), \text{ and} \\ 0, & \text{otherwise.} \end{cases}$$

This ξ is clearly in $M(\mathcal{S}, q)$, and so for every $\pi \in (\mathcal{S}, q)$, $\pi \cdot \xi = 0$ is required. But this is

$$-\pi_{f^0_\tau} q_{sf^0_\tau} + \sum_{f_t \in \mathbf{S}(f^0_\tau)} \pi_{f_t} d_s(f_t) \quad = \quad 0.$$

Solve for $q_{sf^0_\tau}$, and you have the result.

(b) Let y be any legitimate trading plan for the securities in $\mathcal{S}$; the wealth transfer vector that is generated by y is ξ given by

$$\xi(f_t) = \sum_{s\in\mathcal{S}} \left(q_{sf_t}\left[y(\hat{f}_t, s) - y(f_t, s)\right] + y(\hat{f}_t, s)d_s(f_t)\right] \Bigg).$$

Hence $\pi \cdot \xi$ is

$$\pi \cdot \xi = \sum_{f_t}\sum_{s\in\mathcal{S}} \left(\pi_{f_t} q_{sf_t}\left[y(\hat{f}_t, s) - y(f_t, s)\right] + \pi_{f_t} y(\hat{f}_t, s)d_s(f_t)\right] \Bigg).$$

Once we show that this is equal to 0 for the q given in the statement of the proposition, we're done: This will show that $\pi \cdot \xi = 0$ for all $\xi \in M(\mathcal{S}, q)$, which (per Proposition 16.8(b)) is one test of viability of $\mathcal{S}$ and q.

To show that the double sum is 0, we will show that, for each s, the sum (only over f_t) is 0. (This shouldn't surprise you; if y is a legitimate trading plan, then so is the plan of executing only trades in a single s that are given by y.) Now, for a fixed security s,

$$\pi \cdot \xi = \sum_{f_t} \left(\pi_{f_t} q_{sf_t}\left[y(\hat{f}_t, s) - y(f_t, s)\right] + \pi_{f_t} y(\hat{f}_t, s)d_s(f_t)\right] \Bigg).$$

Substitute the formula for q_{sf_t}, and this becomes

$$\sum_{f_t} \left\{ \pi_{f_t} \left(\frac{1}{\pi_{f_t}} \left[\sum_{f'_{t'} \in \mathbf{S}(f_t)} \pi_{f'_{t'}} d_s(f'_{t'}) \right] \right) \left[y(\hat{f}_t, s) - y(f_t, s)\right] + \pi_{f_t} y(\hat{f}_t, s) d_s(f_t) \right\},$$

which is

$$\sum_{f_t} \left\{ \left[\sum_{f'_{t'} \in \mathbf{S}(f_t)} \pi_{f'_{t'}} d_s(f'_{t'}) \right] \left[y(\hat{f}_t, s) - y(f_t, s)\right] + \pi_{f_t} y(\hat{f}_t, s) d_s(f_t) \right\}.$$

In these manipulations, you may be worried about the fact that some f_t may not be legitimate trading dates for t. But for those dates, the term $y(\hat{f}_t, s) - y(f_t, s)$ is zero, so it is irrelevant to the sum if we put these terms in, with what the formula says *would have been* the price of s, had it traded. Now for each f_t, recall that $\mathbf{P}(f_t)$ denotes the set of all predecessors of f_t, and recall that we have invented a predecessor for f_0 at which y is zero. I want to look at the term in the last display, running the sum over f_t in which s pays a nonzero dividend. This is

$$\sum_{f_t} \pi_{f_t} d_s(f_t) \left[y(\hat{f}_t, s) + \sum_{f'_{t'} \in \mathrm{P}(f_t)} \left(y(\hat{f}'_{t'}, s) - y(f'_{t'}, s) \right) \right].$$

Let me translate this: Two displays ago, we had the sum, for each f_t, of the trading plan's change in portfolio times the sum of normalized future dividends (normalized by π for the contingency in which the dividend is paid), plus the holdings of the previous contingency times any normalized dividend paid immediately. In the display immediately above, we look at all terms associated with a dividend paid in f_t, and sum over all the f_t. We have the normalized dividend multiplied by: The holdings from last period, and then the sum of all changes in holdings, going back to $\hat{f}_0$, in all predecessors of f_t. But, if you think about it just for a moment, you'll see that the term in square-brackets in the last display is a telescoping sum that, after cancelling out the same terms with different sign, leaves us with $y(\hat{f}_0, s)$, which is zero. So the last displayed term is zero, completing the proof.

For the nth time, I caution you not to be too impressed by this. It is (once again) simple bookkeeping and nothing more. If the price of a security is the properly "discounted" sum of its future dividends, then when you buy a unit of security, you are "paying" for those future dividends. And then, for as long as you hold the security, it pays some of those dividends; when and if you sell it, (by the rule that gives q) you get back the properly discounted sum of any dividends you didn't wait long enough to receive.

■ 16.6. Suppose that $\mathcal{S}$ is $\mathcal{S}^{\text{SFCC}}$, and prices q are given such that $\mathcal{S}$ and q are viable. Pick any f_t, $t > 0$, and some $\omega \in f_t$. (That is, ω is a state of nature that is still possible in contingency f_t.) If s^ω is the security in $\mathcal{S}^{\text{SFCC}}$ that goes with state ω, then we know that $q_{s^\omega f_t} > 0$ (since s^ω pays a strictly positive dividend in a successor to f_t, namely $\{\omega\}$ at time T). But then the trading plan of buying $1/q_{s^\omega f_t}$ units of s^ω at f_0 and then selling at time t (regardless of the contingency) generates the following wealth-transfer vector: At f_0, you pay $q_{s^\omega f_0}/q_{s^\omega f_t}$. At all other dates, except for t, you get nothing. And at date t, *if* the contingency is not f_t, you get nothing (for all other time-t contingencies, ω has been ruled out, the security is defunct, and so it is worthless), while in contingency f_t, you get 1 unit of numeraire. Hence, this trading plan transfers wealth from f_0 to f_t (and to no other contingency), satisfying the requirements of Lemma 16.10. The dimension of $M(\mathcal{S}^{\text{SFCC}}, q)$ must be $N-1$.

Supppose that $\mathcal{S} = \mathcal{S}^{\text{RFCC}}$ and q are viable. For each f_t, let s^{f_t} be the security that trades at $\hat{f}_t$ and pays (only) at f_t. As before, we create a trading plan that moves wealth from f_0 to f_t, for arbitrary f_t, and to no other contingency. Fixing f_t (for $t > 0$), let the path of predecessors of f_t from f_0 to f_t be denoted $f_0, f_1, \ldots, f_t$. The trading plan begins with the purchase of units of s^{f_1} in f_0. If the next contingency is f_1,

s^{f_1} pays a dividend, and the proceeds are used to buy units of s^{f_2}. (If the time 1 contingency is not f_1, no dividend is paid, and the plan involves no more trading.) Then if f_2 happens, s^{f_2} pays a dividend, used to buy units of s^{f_3}, and so forth. Clearly, as time moves from 0 to t, as long as f_t is still possible, the plan enables the purchase of units of the next element of $\mathcal{S}^{\text{RFCC}}$; if at time t, f_t is the contingency, the units of s^{f_t} that were purchased at time $t-1$, in contingency f_{t-1}, pay a dividend, which is "cashed out." Hence, this plan moves wealth from f_0 to f_t (and to no other contingency). Once we get the scale of this plan right (how many units of s^{f_1} are bought at f_0, to get the ball rolling, so that in the end we wind up with 1 unit of s^{f_t} and hence a dividend of 1 unit of the numeraiare), we have satisfied the requirements of Lemma 16.10, and we know that $M(\mathcal{S}^{\text{RFCC}}, q)$ has dimension $N-1$.

I'll leave the details of the argument here, except to say: To be precise about this, you first have to say that each s^{f_t} has a strictly positive price under q in the one contingency (namely, $\hat{f}_t$) in which it trades. But that is another application of first part of Lemma 16.11.

■ 16.8. Suppose we arbitrarily set $\pi = (1, 1, \ldots, 1)$ and use it and formula (16.4) to compute the prices of the L securities at the various time t-contingencies. The formula says that the price of security s for this π in contingency f_t will be

$$q_{sf_t} = \sum_{\omega \in f_t} d_s(\{\omega\}).$$

Since the f_t partition Ω, this tells us that, with probably one, all the square sub-matrices of the $L \times N_t$ matrix whose component corresponding to security s (which indexes the L dimension) and the time-t contingency f_t (which indexes the N_t dimension) is q_{sf_t} have full rank, with probability 1. Since there are only finitely many dates, this is true for all dates $t = 0, \ldots, T-1$ with probability one, as well as for the original matrix of dividends.

I assert that if this property holds (which it does, with probability 1), then for this q, $M(\mathcal{S}, q)$ has dimension $N-1$. The argument is: Take any f_t, $t < T$, and let f_{t+1}^i enumerate its $\ell(f_t)$ immediate successors. We know that $L \geq \ell(f_t)$, so choose any $\ell(f_t)$ securities and look at the $\ell(f_t) \times \ell(f_t)$ square matrix of the prices of these securities at date $t+1$ in the $\ell(f_t)$ contingencies f_{t+1}^i. (If $t = T-1$, look at the matrix of dividends paid by the $\ell(f_t)$ securities you select.) This has full rank, which means that, for each i, a portfolio of these $\ell(f_t)$ securities can be constructed in contingency f_t that is worth 1 in f_{t+1}^i and 0 otherwise. We know that these portfolios all have positive prices (in fact, given our choice of π, we know that each of these portfolios costs 1 unit exactly!), so these portfolios allow us to transfer wealth from any f_t to any one of its immediate successors, without generating gains or losses in any other contingency. This is enough to know that markets are complete. (This isn't quite the criterion given in Lemma 16.10, but if you look at the solution to the second half of 16.6, given previously, you'll see why it is true.)

But then, by exactly the argument given in the text, if $(p, \boldsymbol{x})$ is a Walrasian equilibrium, then for some $\boldsymbol{y}$ (which needs to be constructed, based on $\boldsymbol{x}$), this specific q, and for $r \equiv p$, $(r, q, \boldsymbol{x}, \boldsymbol{y})$ is an EPPPE.

Compare this with Proposition 16.13. That proposition says that, to get complete markets, we need *at least* $\ell(f_t)$ non-defunct securities in contingency f_t. This is almost a perfect converse; in that it says that, as long as securities trade in every non-terminal contingency and have positive dividends in every terminal contingency, any "random selection" of dividends for $\max_{f_t} \ell(f_t)$ securities will do.

But this trades rather heavily on the fact that we are allowed to choose π as we want. How does this work, if we instead have securities whose dividends are denominated in wheat? In fact, it works almost as well, in the following sense. Fix some Walrasian equilibrium $(p, \boldsymbol{x})$. We know that to get this equilibrium with an EPPPE where the price of wheat in every spot market is 1, we'll need $\pi_{f_t} = p_{1f_t}$. So now, when we look at the prices of the securities, instead of getting the simple sum of dividends, we get

$$q_{sf_t} = \frac{1}{p_{1f_t}} \sum_{\omega \in f_t} p_{1\{\omega\}} d_s(\{\omega\}).$$

The $1/p_{1f_t}$ is irrelevant but, when we take the sum of dividends, we are taking weighted sums. So we need to know: If we have a big matrix with each component randomly selected (in a fashion absolutely continuous with respect to Lebesgue measure), and we take weighted sums of the columns, the matrix that results (for specific weights) has full-rank square submatrices with probability one. This is true, if the weights are fixed in advance. But: the weights change with the Walrasian equilibrium we are looking at. *If* we know that the all-markets-at-once economy has only finitely many Walrasian equilibria, which is true for "most well-behaved" pure exchange economies, then we get the probability one statement for all of them at once. So if this *if* is met, this is almost as good as the result for economies with financial securities.

However: (1) In either case, the result is only true for "most" selections of dividends. It is a probability one statement. $L \geq \max_{f_t} \ell(f_t)$ securities might have their dividends chosen in a way that means markets are not complete. And (2) this says that, with probability 1, every Walrasian-equilibrium consumption allocation is an EPPPE consumption allocation. The converse is not claimed. For more on this, I urge you to try Problem 16.9.

■ 16.11. We know that $\mathcal{S}$ and q are viable if and only if there is some strictly positive π such that $\pi \cdot \xi = 0$ for all $\xi \in M(\mathcal{S}, q)$, and that the dimension of $M(\mathcal{S}, q)$ is $N - 1$ if and only if there is, up to normalization, a unique π of this character.

In the problem, it assumes the existence of one security s^0 that pays a dividend of 1 at every terminal (time-T) contingency and whose price, by choice of numeraire, is 1 in every contingency. In this context, imagine a strictly positive π such that $\pi \cdot \xi = 0$. In particular, look at any such π normalized so that $\pi_{f_0} = 1$.

One $\xi \in M(\mathcal{S}, q)$ is obtained by purchasing a share of the distinguished security in f_0 and holding it until time T, collecting dividends of 1. For this ξ, $\pi \cdot \xi = -\pi_{f_0} q_{s^0 f_0} + \sum_{\omega \in \Omega} \pi_{\{\omega\}} d_{s^0}(\{\omega\})$. Since $q_{s^0 f_0} = d_{s^0}(\{\omega\}) = 1$ for every ω, and since we have normalized π_{f_0} to be 1, this being equal to 0 is the same as

$$1 = \sum_{\omega \in \Omega} \pi_{\{\omega\}}.$$

So if we define, for each ω, $\mu(\omega) = \pi_{\{\omega\}}$, $\mu : \Omega \to R_{++}$ and sums to 1; μ constitutes a probability measure on Ω.

Moreover, another $\xi \in M(\mathcal{S}, q)$ is obtained by purchasing a share of the distinguished security in f_t and holding it until time T, collecting dividends of 1 in all terminal contingencies that are successors to f_t. If $\pi \cdot \xi = 0$ for this ξ, and if the price of s^0 is 1 in contingency f_t (which we've assumed), then by the same argument, we get

$$\pi_{f_t} = \sum_{\omega \in f_t} \pi\{\omega\} = \sum_{\omega \in f_t} \mu(\omega),$$

which is to say that π_{f_t} is the probability of contingency f_t under the probability distribution μ.

Now take any other security s (all securities trade in all contingencies, recall), and take any contingency f_t. Consider the trading rule of buying one unit of s in f_t and, if $t < T - 1$, selling it next period, while if $t = T - 1$, holding it for the dividends it will pay. This generates a ξ, and for $\pi \cdot \xi$ to equal 0 for this ξ, it is necessary that

$$\pi_{f_t} q_{s f_t} = \sum_{f_{t+1} \in \mathsf{S}(f_t)} \pi_{f_{t+1}} q_{s f_{t+1}},$$

where if $t = T - 1$, we interpret $q_{s f_{t+1}}$ as $d_s(f_{t+1})$. But if we replace the components of π with $\mu(f_t)$, the probability of contingency f_t under the probability measure μ, this is

$$q_{s f_t} = \sum_{f_{t+1} \in \mathsf{S}(f_t)} \frac{\mu(f_{t+1})}{\mu(f_t)} q_{s f_{t+1}},$$

with the same special treatment for $t = T - 1$, which says that, at contingency f_t, the conditional expected price next time of security s (or the conditional expected dividend it will pay, in the special case $t = T - 1$), is equal to its current price. That is, μ turns the price process of s into a martingale.

So, if there is some $\pi \in (\mathcal{S}, q)$, which is true if and only if $\mathcal{S}$ and q are viable, then there is an associated probability distribution μ under which every security is a martingale. And, if you run all the arguments we've given backwards (or use formula

(16.4) and, in particular, Lemma 16.11b, the converse emerges: If μ is a full-support probability on Ω that turns each security's price process (with the dividend at time T) into a martingale, then π_{f_t} equal to the probability of f_t (under that probability) is a strictly positive element of R^N that satisfies $\pi \cdot \xi = 0$ for every $\xi \in M(\mathcal{S}, q)$. So $\mathcal{S}$ *and* q *are viable (where* $\mathcal{S}$ *conforms to the assumptions of this problem) if and only if there is a (strictly positive) probability* μ *on* Ω *that makes each security price process into a martingale.*

And, moreover, for each $\pi \in (\mathcal{S}, q)$ normalized so that $\pi_{f_0} = 1$, there is a different "martingale measure" μ on Ω. So *the dimension of* $M(\mathcal{S}, q) = N - 1$, *which is to say that markets are complete, if and only if there exists a unique "martingale measure" for the given data.*

Although you weren't asked to show this, we can go even a step further: Suppose you are given a set of securities $\mathcal{S}$ with the properties of this problem and prices q for those securities, and you are asked: Suppose we add another security s^* to the set, which trades in every contingency and pays dividends only at date T. Can we put any bounds on what the equilibrium price of s^* might be, assuming that adding it leaves $\mathcal{S}$ and q unchanged? The answer is, The price of s^* in any contingency f_t must lie within the range of conditional expected values of its dividends, where we take conditional expectations with respect to all the "martingale measures" for $\mathcal{S}$ and q. In particular, if all those conditional expectations are the same, then s^* is "priced by arbitrage," meaning there is a trading plan involving the securities in $\mathcal{S}$ that precisely replicates what s^* does.

■ 16.13. To give a version of Proposition 16.15 that includes firms, we need to revise definitions a bit. We continue to suppose that there is a subspace $\mathcal{M}$ of R^{kN} of "traded bundles," and prices will be given by a linear functional $\pi : \mathcal{M} \rightarrow R$. Consumers will look at net trades from $\mathcal{M}$, and they will be assumed to be locally insatiable constrained to $\mathcal{M}$, as before. But we add: Firm f is characterized by a production-possibility set $Z^f \subseteq \mathcal{M}$ and by the desire, facing prices given by π, to maximize profit, given by $\pi(z)$ for $z \in Z^f$. Consumer h holds an s^{fh} share in firm f, where $\sum_h s^{fh} = 1$ for each f, and all the shares are nonnegative. An $\mathcal{M}$-constrained Walrasian equilibrium is a linear functional $\pi : \mathcal{M} \rightarrow R$, net trades ζ^h for each consumer, and production plans z^f for each firm, such that $\zeta^h \in \mathcal{M} \cap \mathcal{X}^h$ and ζ^h maximizes u^h over all $\hat{\zeta}^h$ that satisfy h's budget constraint $\pi(\hat{\zeta}^h) \leq \sum_f s^{fh}\pi(z^f)$, z^f maximizes $\pi(\hat{z}^f)$ over all $\hat{z}^f \in Z^f$, and markets clear: $\sum_h \zeta^h \leq \sum_f z^f$.

And then, in this expanded setting, we want to prove that if $(\pi, (\zeta^h)_h, (z^f)_f)$ is an $\mathcal{M}$-constrained Walrasian equilibrium for $\mathcal{M}$-constrained locally insatiable consumers and for π a nonnegative linear functional on $\mathcal{M}$, then no feasible $\mathcal{M}$-constrained consumption plan $(\hat{\zeta}^h)_h$ can Pareto-dominate (ζ^h), where feasibility means that $\hat{\zeta}^h \in \mathcal{M} \cup \mathcal{X}^h$ for each h, and there are feasible production plans $\hat{z}^f \in Z^f$, one for each firm, such that $\sum_h \hat{\zeta}^h \leq \sum_f \hat{z}^f$.

To prove this, suppose that $(\hat{\zeta}^h)$ is feasible, and the plans $(\hat{z}^f)$ are what make it so.

Then since each $\hat{\zeta}^h \in \mathcal{M}$ and each $\hat{z}^f \in \mathcal{M}$, so is $\sum_h \hat{\zeta}^h - \sum_f \hat{z}^f$. Feasibility ensures that $\sum_h \hat{\zeta}^h - \sum_f \hat{z}^f \leq 0$. So, since π is nonneative,

$$0 \geq \pi\left(\sum_h \hat{\zeta}^h - \sum_f \hat{z}^f\right) = \sum_h \pi(\hat{\zeta}^f) - \sum_f \pi(\hat{z}^f), \quad \text{or} \quad \sum_h \pi(\hat{\zeta}^h) \leq \sum_f \pi(\hat{z}^f).$$

Of course, since z^f maximizes firm z's profit at the prices π, $\pi(\hat{z}^f) \leq \pi(z^f)$ for each f, and so the last inequality in the display implies

$$\sum_h \pi(\hat{\zeta}^h) \leq \sum_f \pi(z^f). \tag{$\star$}$$

Now suppose that $(\zeta^h)_h$ is Pareto dominated by $(\hat{\zeta}^h)$. Since each consumer h is $\mathcal{M}$-constrained locally insatiable, we know that each h spends all her wealth in equilibrium. (If she didn't, she could find some other $\zeta' \in \mathcal{M} \cup \mathcal{X}$ arbitrarily close to ζ^h—close enough to still obey her budget constraint, which is strictly better for her than ζ^h.) So, for each h, $\pi(\zeta^h) = \sum_f s^{fh}\pi(z^f)$. And if $(\hat{\zeta}^h)_h$ Pareto dominates (ζ^h), then (1) by local insatiability again, we know that $\pi(\hat{\zeta}^h) \geq \pi(\zeta^h)$ for each h, and (2) for any h (and there must be one) such that $\hat{\zeta}^h$ is strictly preferred to ζ^h, $\pi(\hat{\zeta}^h) > \pi(\zeta^h)$. So summing over h,

$$\sum_h \pi(\hat{\zeta}^h) > \sum_h \pi(\zeta^h) = \sum_h \sum_f s^{fh}\pi(z^f) = \sum_f \pi(z^f)\left[\sum_h s^{fh}\right] = \sum_f \pi(z^f).$$

This directly contradicts ($\star$); a feasible $(\hat{\zeta}^h)_h$ cannot Pareto-dominate $(\zeta^h)_h$.

■ 16.15. We are imagining an economy with consumers and firms. The description of consumers doesn't change from before, except that they will be endowed with initial shareholdings to be described later. There are also non-equity securities $\mathcal{S}$, just as before. For simplicity, I'll assume that these securities are financial; they pay dividends in units of account or numeraire, although it isn't hard to adapt this formulation to a case where all dividends (including those paid by firms) are paid in some single commodity whose equilibrium price is ensured to always be positive and that is chosen to be the numeraire in each contingency. The symbols x^h and $\boldsymbol{x}^h$ for consumption plans for consumer h, y^h and $\boldsymbol{y}^h$ for trading plans (in securities in h) for consumer h, r for prices in the spot markets, and q for prices for the securities in $\mathcal{S}$ continue to be used.

Firms are described by production-possibility sets $Z^f \subseteq R^{kN}$. (The double use of f for firms and for contingencies can't be avoided, unless we use a different letter for one or the other. I think it is less confusing to just go with the double use, so I do that.) If a firm chooses production plan z^f, it must *finance* this plan, which means it must issue dividends $d_f(f_t)$ and have a financial plan which involves holding a portfolio of the securities in $\mathcal{S}$ in each contingency f_t, where $w^f(f_t, s)$ will denote firm f's (post-) contingency f_t holdings of security s for $s \in \mathcal{S}$. Firm f's production, dividends, and financing plan (z^f, d_f, w^f) must satisfy two constraints: The firm's financial plan must respect any trading constraints on the securities $s \in \mathcal{S}$. And the firm must *balance its books* in every contigency, meaning that, given prices r and q, for each f_t,

$$r_{f_t} \cdot z^f_{f_t} + \sum_{s \in S} w^f(\hat{f}_t, s) d_s(f_t) = d_f(s_t) + \sum_{s \in S} q_{sf_t} \left[w^f(f_t, s) - w^f(\hat{f}_t, s) \right].$$

To translate: $r_{f_t} \cdot z^f_{f_t}$ is the net operating income of the firm, the market value of its outputs less the cost of inputs. Added to this are any dividends the firm might receive from the portfolio of securities it holds; of course, this is based on its holdings the contingency before, because we assume that securities trade *ex dividend*. On the right-hand side is the dividend this firm pays, plus the net cost to it of transactions it undertakes in the securities markets. Suppose, for instance, the firm is just starting out at f_0. It begins with no security holdings. And suppose that, at f_0, it must purchase inputs, so $r \cdot z < 0$. Then to achieve an equality in the "balance the books" equation for f_0, the firm must sell some securities, which, in real life, is called "borrowing some money" or "issuing debt," although as we noted in the text, without bankruptcy, debt is debt. A richer formulation could have it *float equity*, but I'll avoid that sort of thing in this answer. If you are courageous, add it in. Also, I am not allowing firms to trade in each other's equity or, even worse, to buy and sell its own equity or create securities not in $\mathcal{S}$. So long as markets are complete using $\mathcal{S}$ alone—which will be assumed—any security the firm created would not add trading possibilities, hence would be economically redundant. Of course, in real life, these things happen, and you can modify my formulation so that this is allowed. It is tedious but doable. But if $\mathcal{S}$ doesn't provide complete markets on its own, the difficulties you will encounter will be immensely greater.

We now add equity in each firm to the story. For each firm f, there is one more security added to the securities markets, equity in the firm. I'll use the symbol σ^f to represent a "share" of equity in the firm, where we normalize things so that there is only one share in the firm in existence. I'll assume that equity in a firm trades in every contingency f_t for $t < T$. Dividends are determined endogenously by firms, as part of their production, dividend, and financing plans. Only consumers (in this formulation) are allowed to hold equity in firms: A trading plan y^h has components $y^h(f_t, s)$ for $s \in S$ and $y^h(f_t, \sigma^f)$ for $f \in F$. Recall that $\hat{f}_0$ is used conventionally to describe the "predecessor" to f_0 for budget-equation purposes; when we write (say) $y^h(\hat{f}_0, s)$, we always mean 0, since this represents the endowment of consumer holdings of securities in $\mathcal{S}$. Now we continue to use $y^h(\hat{f}_0, \sigma^f)$, but these numbers are exogenously given, with the stipulation that they are always nonnegative and sum (over all consumers, for each firm) to 1. The consumer's budget constraints are changed to reflect these new securities, where $q_{\sigma^f f_t}$ is the equilibrium price of the (one) share of firm f's equity. (I reiterate, if you want to change my formulation to allow firms to float their own equity, go at it. But if you do this, since the value of a firm's equity will not always be zero, you will need to have "founders" or "entrepreneurs" in your model for each firm, who reap the benefits of the initial public offering. The nice thing is, the meta-theorem will still apply, as long as its rather strong assumptions are maintained.)

In this formulation, you must be careful about dividends paid by the firm. Since firms have shareholders from the start (at least, the way I've done this), you can allow the firm to pay dividends even at f_0, something we didn't allow for $s \in \mathcal{S}$. Also, the issue of negative dividends must be addressed. You can restrict the firm to pay only non-negative dividends. But if you do, you should probably assume that $0 \in Z^f$ for each firm, so you know (once you define profit, see below) that the firm can achieve zero profit. And this may make some level of participation in the securities markets a requirement, at least for some production plans. Imagine, for instance, a very productive z^f that, unhappily, requires significant inputs at time T to "clean up" (a legal requirement, say) after producing tons and tons of really valuable output at time $T-1$. To make this feasible with nonnegative dividends, the firm had better buy some securities at time $T-1$ that will allow it to finance the required operations at time T.

That completes the definition of the economy. (Whew!) The next step is to define an EPPPE. This, it turns out, is where things get murky. The problem, as described informally in the text, is: What is the profit of the firm? We need to have an objective function for the firm, but what is it?

To get to a formal statement of the meta-proposition, we duck this issue: We only define EPPPE with complete markets:

Definition G16.1. *A **complete-markets EPPPE** is an array* $(r, q, \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}, \boldsymbol{d}, \boldsymbol{w})$, *consisting of:*

a. *consumption and trading plans* $\boldsymbol{x}^h$ *and* $\boldsymbol{y}^h$ *for each consumer (where the trading plans include the possibility of trading in the equity of firms,*

b. *production, dividend, and financing plans* $\boldsymbol{z}^f$, $\boldsymbol{d}_f$, *and* $\boldsymbol{w}^f$ *for each firm,*

c. *spot-market prices* r_{f_t} *for the commodities in each* f_t,

d. *prices* q_{sf_t} *for the non-equity securities in* $\mathcal{S}$, *and*

e. *prices* $q_{\sigma^f f_t}$ *for the equity of each firm* f *in each contingency* f_t,

such that

f. *the plans (consumption and trading) of each consumer respect the rules of trading in securities in* $\mathcal{S}$ *and satisfy the consumer's budget constraints,*

g. *the plans (production, dividend, and financial) of each firm respect the rules of trading in securities in* $\mathcal{S}$ *and satisfy the firm's "balance the books" constraints in each* f_t*,*

h. *each consumer* h *maximizes her utility with* $\boldsymbol{x}^h$*, in comparison with any other* $\hat{x}^h$ *for which there is a trading plan* $\hat{y}^h$ *that meets all the constraints of part f,*

i. *the dimension of* $M(\mathcal{S}, q)$ *is* $N-1$*, (with* $\pi \in R^N_{++}$ *the unique vector in* $(\mathcal{S}, q)$ *normalized so that* $\pi_{f_0} = 1$*),*

j. *each firm* f *maximizes the value of its initial equity,* $q_{\sigma^f f_0} + \boldsymbol{d}_f(f_0)$*, as determined by its production, dividend, and financing plan* $\boldsymbol{z}^f$*,* $\boldsymbol{d}_f$*, and* $\boldsymbol{w}^f$*, in comparison with any other set of plans* $\hat{z}^f$*,* $\hat{d}_f$*, and* $\hat{w}^f$ *that together satisfy the constraints imposed in part g, where the firm asseses the value of its initial equity under the alternative, hat plan, as*

$$\sum_{f_t \in \mathcal{F}} \hat{d}_f(f_t)\pi_{f_t}, \quad \textit{and}$$

k. *all markets clear: with inequalities in the spot-commodity markets, with equality (for a net supply of zero) in the markets for securities in* $\mathcal{S}$*, and with equality (for a net supply of 1) in the markets for each firm's equity.*

The key to this definition is part j, which provides the objective function for the firm: it seeks to maximize the *initial value of its equity*, taking all prices, and in particular the implicit pricing vector π, as given. It may not be one-hundred percent clear that $q_{\sigma^f f_0} + d_f(f_0)$ is equal to $\sum_{f_t \in \mathcal{F}} \boldsymbol{d}_f(f_t)\pi_{f_t}$, so this needs to proved. (On the other hand, if you understood the proof of Lemma 16.11, this may be very clear.) It should be noted that the term *the value of [firm* f*'s] initial equity* we mean the price the equity commands in the f_0 market plus any dividend the firm might choose to pay to its initial shareholders at f_0. (Securities are priced *ex dividend*, so the value of this dividend is not impounded in $q_{\sigma^f f_0}$.) Why did we choose this as the objective of the firm? We have the following result:

Proposition G16.2. *Suppose that* $(r, q, \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}, \boldsymbol{d}, \boldsymbol{w})$ *is a complete-markets EPPPE.*

a. *For each firm* f*,*

$$\boldsymbol{d}_f(f_0) + q_{\sigma^f f_t} = \sum_{f_t \in \mathcal{F}} \pi_{f_t} r_{f_t} \cdot \boldsymbol{z}^f_{f_t}.$$

b. For each f_t*,*

$$q_{\sigma^f f_t} = \frac{1}{\pi_{f_t}} \sum_{f'_{t'} \in S(f_t)} \pi_{f'_{t'}} d_f(f'_{t'}).$$

c. Suppose that any firm f *contemplates shifting its dividend and financing structure from* $\boldsymbol{d}_f$ *and* $\boldsymbol{w}^f$ *to some alternative plan* $\hat{d}_f$ *and* $\hat{w}^f$ *that, with* $\boldsymbol{z}^f$*, continues to satisfy the constraints in part j of the definition. Then* $(r, \hat{q}, \boldsymbol{x}, \hat{y}, \boldsymbol{z}, \hat{d}, \hat{w})$ *is a complete-markets EPPPE, where* $\hat{q}$ *is the same as the initial* q*, except for some possible repricing of the equity of this firm* f*,* $\hat{y}$ *involves adjusting the trading plans of consumers to accomodate the net dividend structure of firm* f *and its new financing plan,* $\hat{d}$ *is the same as* $\boldsymbol{d}$ *except for the one firm* f*, which shifts to* $\hat{d}^f$*, and* $\hat{w}$ *is the same as* $\boldsymbol{w}$*, except that firm* f *shifts to* $\hat{w}^f$*.*

To explain: Part a says that the value of initial equity is the discounted sum of contingency-by-contingency net operating profits of the firm. So, when we said that the objective of the firm is to maximize the value of its initial equity, we could just as well have defined its "profit" as the discounted sum of these net operating profits (where the discount factors π_{f_t} are driven by equilibrium prices for the securities in $\mathcal{S}$) and said that this is what firms maximize. Part b confirms that the formula used by the firm to evaluate alternatives to the plan it chooses is the formula that determines the value of its initial equity. As for part c, that is the formal statement of the famous Modgliani–Miller Theorem—-that the capital structure of a firm (what combination of debt and equity it uses to finance its operations) is irrelevant in perfect (or frictionless) markets, because investor/consumers can always use other markets to "undo" what the firm does in setting up its capital structure. Of course, the *capital structure* of a firm will affect the price of its equity (although *not* the value of its initial equity—that's what part a tells us); a firm that pays a big dividend in some contingency will have a lower value of equity (*ex dividend*) than a firm that takes the money that would have been used to pay that dividend and invests it for a while in securities from $\mathcal{S}$.

The proof of this proposition is one more exercise in accounting and arbitrage arguments. For part a, you have to construct an argument somewhat similar to the proof of Lemma 16.11b, where you show that the operating profit accrued in some contingency f_t has to reach shareholders (at a market-driven, fair rate of return) at some point in time. Part b follows the proof of Lemma 16.11a precisely. As for part c, the only difficult part is to show that, if one firm changes its dividend and financing plan, consumers can take the opposite side of its transactions without affecting the wealth they have in each contingency to purchase real commodities.

And, with those hints, I leave the details to you. Next comes

Proposition G16.3. *Suppose that* $(r, q, \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}, \boldsymbol{d}, \boldsymbol{w})$ *is a complete-markets EPPPE. Then there is a Walrasian equilibrium* $(p, \boldsymbol{x}, \boldsymbol{z})$ *for the all-markets-at-once economy (where shareholdings in the latter are the initial shareholdings in the dynamic economy). And, conversely, if*

$(p, \boldsymbol{x}, \boldsymbol{z})$ is a Walrasian equilibrium for the all-markets-at-once economy and if $\mathcal{S}$ is rich enough so that $M(\mathcal{S}, q)$ has dimension $N-1$, where q is defined for securities in S by formula (16.4), for some strictly positive $\pi \in R^N_{++}$, then there is a complete-markets EPPPE in which firms produce according to $\boldsymbol{z}$ and consumers consume according to $\boldsymbol{x}$.

This is another exercise in bookkeeping, which I leave for you.

There is one final point to address: In the text, I said something about how, in this sort of complete-markets EPPPE, shareholders unanimously prefer the production plan chosen by the firm over any other feasible plan (taking prices as given), and—subject to a mysterious caveat—this continues to be true of subsequent shareholders.

I can finally explain the mysterious caveat. In contingency f_t (say), the firm will be committed to having done certain things. It will have undertaken certain production activities. It will be holding a portfolio of securities and have paid dividends. None of this can be undone, and when I say that current shareholders prefer that the firm carry out its plans going forward, I mean, taking all those past actions as given; no time machines are available to undo what has been done. But if we allow any specification of $Z^f \subseteq R^{kN}$, the possibility exists that what the firm is capable of depends on what it would have done in contingencies now known to be impossible. Go back to the discussion in the text of the consumer whose marginal utility for consuming wheat in one contingency depends on what she would have consumed in another contingency that is now known not to be happening; this is just the "production" equivalent. For production technologies of this sort, shareholders at time t don't need a time-machine to say things like, "The firm planned to undertake $z^f_{f'_t}$ if today's contingency was f'_t instead of f_t, perhaps because this would have generated big operating profits and, hence, dividends. But we know f'_t isn't going to happen, and if the firm shifts what it would have done, Z^f is constructed so that it can do more (make more operating profits) today and going forward. So we like that shift."

This phenomenon, which is possible in theory, will prevent the "shareholders continue to unanimously prefer the original plan" from being true. What is needed is a restriction on feasible production technologies that rules this out; essentially, that says that things the firm might have done in a different state of nature doesn't constrain what it can do today and going forward. Only the past that was constrains the firm going forward; not the present and future that will never be. Formalized, this means some "state-by-state separability" in the construction of Z^f; if that separability is present, then the "shareholders continue to prefer…" result can be obtained. (You are invited to do so.)