Microeconomic Foundations I:
Choice and Competitive Markets

Student's Guide

Chapter 8: Social Choice and Efficiency

Summary of the Chapter

Social Choice Theory concerns the following problem. A set of individuals or households H and a set of social states X are given. Each individual has her own preferences concerning the various social states, and we want to *aggregate* those preferences, with a view towards choosing a social state from all of X, if all the social states are feasible, or perhaps from some subset of feasible states $A \subseteq X$. Moreover, we want to do this in a fashion that properly reflects the preferences of the individuals, where the definition of the term "properly" is part of the theory.

The chapter provides two answers to the question, How is this to be done? Sections 8.1 and 8.2 concern *Arrow's Theorem*, which says, in essence, that no completely satisfactory solution to the problem is possible. Assume that

1. The set of social states X is finite
2. The set of individuals H is finite
3. The preferences of individual h are given by a complete and transitive preference relation $\succeq_h$ defined on X

We look for a social preference function Φ, which maps every H-tuple of preferences $(\succeq_h)_{h\in H}$ into a complete and transitive preference relation $\succeq = \Phi[(\succeq_h)_{h\in H}]$. The assumption that Φ should work on every H-tuple of complete and transitive preferences is called the *universal domain* assumption. The assumption that the value of Φ at each argument is a complete and transitive preference relation on X is called the *coherence* assumption. These assumptions on the domain and range of Φ are stated formally in the text as Assumption 8.1.

We employ short-hand notation $\succeq$ for $\Phi[(\succeq_h)_{h\in H}]$ and $\succeq'$ for $\Phi[(\succeq'_h)_{h\in H}]$. The strict preference relation derived from $\succeq_h$ is denoted $\succ_h$; $\succ$ is strict preferences derived from $\succeq$, and so forth.

We don't want just any social preference function Φ, but one that has desirable properties and avoids undesirable properties. Definition 8.2 gives two properties that are deemed to be desirable and then a third, undesirable property:

Definition 8.2.

a. *The social preference function Φ satisfies **unanimity** if, for any profile of individual preferences $(\succeq_h)_{h\in H}$ and any pair of social states x and y, if $x \succ_h y$ for each $h \in H$, then $x \succ y$.*

b. *The social preference function Φ satisfies **independence of irrelevant alternatives (IIA)** if, for any two profiles of individual preferences $(\succeq_h)_{h\in H}$ and $(\succeq'_h)_{h\in H}$ and any two social states x and y such that $x \succeq_h y$ if and only if $x \succeq'_h y$ for all $h \in H$, $x \succeq y$ if and only if $x \succeq' y$.*

c. *The social preference function Φ is **dictatorial** if there is some $h^* \in H$ such that, for every profile of individual preferences $(\succeq_h)_{h\in H}$ and every pair of social states x and y, $x \succ_{h^*} y$ implies $x \succ y$.*

The text interprets these properties and gives arguments in favor of a and b and against c. Then:

Proposition 8.3. Arrow's Theorem. *Suppose that X contains three or more elements. If Φ satisfies Assumption 8.1 (the universal domain and coherence assumptions,) unanimity, and IIA, then Φ is dictatorial.*

This result tells us that standard social preference functions must be defective in one way or another. Two examples, majority rule and the *Borda Rule* (how athletic teams are often rank-ordered by the votes of sportswriters or coaches), are provided; majority rule can fail to provide transitive social preferences, while the Borda Rule will fail IIA.

Arrow's Theorem tells us that if we want to avoid dictatorial social preference functions, we must give up one (or more) of (a) universal domain, (b) coherence, (c) unanimity, or (d) IIA. Section 8.2 explores three of these options:

- No one seems very interested in giving up unanimity. If every member of society thinks x is strictly better than y, it is hard to imagine that society should conclude otherwise, for any social preference function that is "nice."
- A substantial literature concerns giving up on the universal domain assumption. Representative of and foremost in this literature is the so-called *assumption of single-peaked preferences*. This assumption and its consequences are developed.
- We might think of giving up on IIA. A primary rationalization for IIA is that we cannot make interpersonal comparisons of utility if our data consist only of the ordinal preferences of the individuals. But if we imagine that we have data that

allow us to calibrate intensity of preferences, for example via calibration with lotteries against benchmark social states, then we could use those data to make judgments based on interpersonal comparisons.

But none of these alternatives is viewed as altogether satisfactory, and we move in Section 8.3 to the "solution" adopted by mainstream economics and economists: Give up on producing a coherent (complete and transitive) social ordering. Instead, rank social states by efficiency or the so-called *Pareto ordering*, which is transitive but not complete:

Definition 8.6. *The outcome or social state x is **Pareto superior to** y (or Pareto dominates y), if $x \succeq_h y$ for every h in H and $x \succ_h y$ for at least one $h \in H$. The social state x is **strictly Pareto superior to** y (or strictly Pareto dominates y) if $x \succ_h y$ for all h. For a subset A of X and a point $x \in A$, x is **Pareto efficient** (or just efficient) within A if there is no $y \in A$ that Pareto dominates x. The set of Pareto efficient points within A is called the **Pareto frontier** of A.*

Discussion of the formal properties of Pareto efficiency follows this definition, together with analysis of the question, In applied settings, how would one "compute" the Pareto frontier? The answer to this question that is developed is, Under appropriate convexity assumptions, the Pareto frontier is, more or less, the set of solutions to the problems, Maximize weighted sums of the utilities of the various consumers, maximizing over the possible social states. After formalizing and proving this result, it is applied in Section 8.5 to the problem of *Syndicate Theory*: How should a finite collection of risk-averse expected-utility maximizers efficiently share in a collective set of risks that they face.

The analysis in this Section 8.4 is important for a reason beyond the results given: It is the first application in the text of the Separating-Hyperplane Theorem, which in the rest of the text is far and away the most useful and used hammer in our mathematical toolbox.

The chapter concludes (Section 8.6) with two cautions. Mainstream economics is, for the most part, passionate about the merits of efficiency. Perhaps because economists are generally unwilling to make value judgments among different efficient social states (because they are unwilling to make interpersonal utility tradeoffs), efficiency seemingly becomes the sole virtue.

- But no argument is offered in favor of the proposition that *any* efficient social state is better than any other state that is not efficient. A social state may be efficient and, at the same time, be highly inequitable, on which grounds it may be judged inferior as an outcome to some inefficient outcome. You should guard against choosing one process or policy or mechanism instead of a second merely because the first produces a Pareto-efficient outcome, while the second does not.

- Efficiency is based on the concept of consumer sovereignty. But, especially in dynamic contexts, where tastes change or the population of individuals changes with changes in the social state, applying consumer sovereignty can be problematic. Consumer sovereignty should command great deference and respect, but it should also be approached with at least a modicum of skepticism.

In terms of the problems at the end of this chapter, I urge every reader to try his or her hand at Problem 8.11.

Solutions to Starred Problems

▪ 8.3. (a) The argument that majority rule produces a complete binary relation (for every x and y, either $x \succeq y$ or $y \succeq x$), is established in the text: Since each $\succeq_h$ is complete, either $x \succeq_h y$ or $y \succeq_h x$ for each x. Therefore, at least half the h "vote" one way or the other (and more than half may "vote" each way, in which case $x \succeq y \succeq x$ would be the result of majority rule).

The hard part of the proof is to show that $\succeq$ is transitive. So suppose that $x \succeq y$ and $y \succeq z$, both according to majority rule. We must show that $x \succeq z$.

If $x = y$, then $y \succeq z$ tells us $x \succeq z$. If $y = z$, $x \succeq y$ tells us $x \succeq z$. And if $x = z$, then $x \succeq z$ follows because everyone votes for x being at least as good at itself (which is z). Therefore, we can proceed under the assumption that x, y, and z are distinct.

The proposition (to keep matters simple) assumes that each $\succeq_h$ is anti-symmetric. Since x, y, and z are distinct, for each h and every pair, $\succ_h$ holds between the two.

Now to enlist single-peakedness. Because the number of individuals is odd, we know that at least one individual must hold that $x \succ_h y \succ_h z$, since more than half the h have $x \succ_h y$ and more than half have $y \succ_h z$. By single-peakedness, this implies that z cannot lie between x and y (when these are viewed as numbers). Now ennumerate all the possible cases; each h must belong to one (and only one) of the following six sets:

$$A = \{h \in H : x \succ_h y \succ_h z\}, \qquad B = \{h \in H : x \succ_h z \succ_h y\},$$

$$C = \{h \in H : y \succ_h x \succ_h z\}, \qquad D = \{h \in H : y \succ_h z \succ_h x\},$$

$$E = \{h \in H : z \succ_h x \succ_h y\}, \text{ and } F = \{h \in H : z \succ_h y \succ_h x\}.$$

We know that half or more of the individuals belong to the union of A, B, and E, since these are the three sets with $x \succ_h y$. And since a majority have $y \succ_h z$, half or more must belong to the union of A, C, and D. To show that $x \succeq z$ by majority rule, we must show that half or more belong to the union of A, B, and C.

I assert that either D or E must be empty. We already know (from single-peakedness and the nonemptiness of A) that z cannot lie between x and y. So either x lies between y and z, which implies that D must be empty, or y must lie between x and z, which implies that E is empty. And if D is empty, then A and C are more than half of H, hence $x \succeq z$ by majority rule, while if E is empty, then A and B are more than half of H, and again $x \succeq z$ by majority rule.

(b) Suppose X consists of three (distinct) elements, $X = \{x, y, z\}$. Suppose there are seven individuals. Two have $x \succ_h y \succ_h z$, two have $z \succ_h x \succ_h y$, and three have $y \succ_h x \succ_h z$. Are these preferences consistent with single-peaked preferences? Yes: Suppose $y > x > z$. It is easy to verify that the preferences given are consistent with this geometric ordering and single-peakedness.

In a pairwise comparison of x and y, four out of the seven prefer x, so $x \succ y$. In a pairwise comparison of y with z, five vote for y, so $y \succ z$. (We know from part a, then, that a pairwise comparison of x with z must give a majority in favor of x, and in fact five vote this way.) But if each individual voted for his or her most preferred outcome, x would get two votes, z would get two, and y would get three. If this were a parliamentary election in Great Britain (and many other places), y would be the winner!

■ 8.5. This is an exercise in creative function construction. Let $v^0 = u(x^0)$, and let W be defined by

$$W(v) = \begin{cases} \sum_h (v_h - v_h^0), & \text{if } v \geq v^0, \text{ and} \\ \sum_h -e^{-(v_h - v_h^0)}, & \text{if } v \not\geq v^0. \end{cases}$$

Note that this function satisfies $W(v^0) = 0$ and $W(v) < 0$ for $v \not\geq v^0$. Therefore, it certainly will be maximized at v^0, if we are looking at the maximum over any set that contains v^0 but contains no other point that is $\geq v^0$. It is clearly strictly increasing on its two "pieces" (that is, on $\{v : v \geq v^0\}$ and on $\{v : v \not\geq v^0\}$ separately, as it is the component-by-component sum of strictly increasing functions on each component. And if $v \geq v'$, "across" the two pieces, then it must be that v is in the $\geq v^0$ piece: if $v \geq v' \geq v^0$, then $v \geq v^0$, and both v and v' are in the same piece. But if $v \geq v^0$ and $v' \not\geq v^0$, then $W(v) \geq 0$ while $W(v') < 0$, hence we have $W(v) > W(v')$. (Not only is W strictly increasing across the two pieces, but every $v \geq v^0$ has $W(v)$ strictly greater than every v' such that $v' \not\geq v^0$.)

Needless to say, this is a function that is absolutely tailor-made to pick out v^0, if it is at all feasible to do so!

■ 8.6. Suppose $X = \{x \in R^3 : x \geq 0, x_1 + x_2 + x_3 = 1\}$. That is, X is a two-dimensional unit simplex. H will have two members. The first member of H, denoted 1, has utility function $u_1(x) = x_1^{1/2} + 0.5x_2^{1/2}$. And individual 2 has utility function $u_2(x) = 0.5x_2^{1/2} + x_3^{1/2}$. To see what sort of set of utility imputations this produces, I plotted the utility imputations in Excel for values of (x_1, x_2, x_3) where each coordinate is a multiple of 0.05. The result is shown in Figure G8.1. The graininess of the plot is due, of course, to the graininess in the values of the coordinates of x. I have drawn in (freehand) what would be (approximately) the boundaries of the full set of utility imputations;

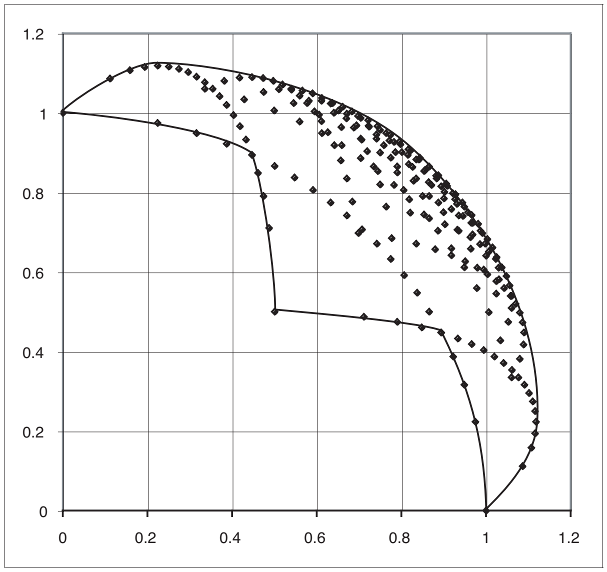

Figure G8.1. Problem 8.6: A Set of Utility Imputations

the set is convex to the north-east (as it must be per Proposition 8.10) but is not at all convex to the south-west.

■ 8.7. Suppose $x = (y_{hs})$ and $x' = (y'_{hs})$ are both feasible. We must show that for any $\alpha \in [0, 1]$, $\alpha x + (1 - \alpha)x'$ is feasible. Note that the sharing rule $\alpha x + (1 - \alpha)x'$ involves sharing rules $(\alpha y_{hs} + (1 - \alpha)y'_{hs})$. Feasibility is a matter of satisfying whichever constraints are imposed:

The adding-up constraint is always imposed: $\sum_h y_{hs} \leq \sum_j z_{js}$. If x and x' both satisfy this constraint, then $\sum_h y_{hs} \leq \sum_j z_{js}$, hence $\alpha \sum_h y_{hs} = \sum_h \alpha y_{hs} \leq \alpha \sum_j z_{js}$, and $\sum_h y'_{hs} \leq \sum_j z_{js}$, hence (1-$\alpha$) $\sum_h y'_{hs} = \sum_h (1 - \alpha)y'_{hs} \leq (1 - \alpha) \sum_j z_{js}$, therefore

$$\sum_h \alpha y_{hs} + \sum_h (1 - \alpha)y'_{hs} = \sum_h [\alpha y_{hs} + (1 - \alpha)y'_{hs}] \leq \alpha \sum_j z_{js} + (1 - \alpha) \sum_j z_{js} = \sum_j z_{js},$$

and $\alpha x + (1 - \alpha)x'$ satisfies the adding-up constraint.

If either the constraint $y_{hs} \geq 0$ or $y_{hs} + e_{hs} \geq 0$ is imposed, and if x and x' satisfy the constraint(s), then it is easy to see (by an argument similar to the one in the previous paragraph) that $\alpha x + (1 - \alpha)x'$ satisfies the constraint(s).

There is finally the possible constraint $\sum_s \pi_h(s)U_h(y_{hs} + e_{hs}) \geq \sum_s \pi_h(s)U(e_{hs})$. Once we show that the functions

$$x \to u_h(x) = \sum_s \pi_h(s)U_h(y_{hs} + e_{hs})$$

are concave in $x = (y_{hs})$ for each h, satisfaction of this constraint by convex combinations of sharing rules that satisfy the constraint is immediate.

So it remains to show that the functions $x \to u_h(x)$ as defined above are concave. You can probably just cite Proposition A3.17(b) as justification for this and satisfy most graders/referees, but in case you want a bit of detail: Fixing x, x', and α, note first that

$$\begin{aligned} U_h([\alpha y_{hs} + (1-\alpha)y'_{hs}] + e_{hs}) =& U_h(\alpha[y_{hs} + e_{hs}] + (1-\alpha)[y'_{hs} + e_{hs}]) \geq \\ & \alpha U_h(y_{hs} + e_{hs}) + (1-\alpha)U_h(y'_{hs} + e_{hs}) \end{aligned},$$

by the concavity of each U_h, therefore

$$x \to U_h(y_{hs} + e_{hs})$$

is concave in x. But then so is $x \to \pi_h(s)U_h(y_{hs} + e_{hs})$, since a positive constant times a concave function is concave, and then so is $u_h(x)$, since the sum of concave functions is concave.

■ 8.10. To begin, let me resummarize what we know. Pareto-efficient sharing rules are precisely the solutions (if any) to problems of the form

$$\text{Maximize } \sum_h \left[\alpha_h \sum_s \pi_h(s)U_h(x_{hs})\right], \quad \text{subject to } \sum_h x_{hs} \leq W_s,$$

plus (perhaps) nonnegativity constraints on the x_{hs}. The maximization problems can be solved state-by-state, per Proposition 8.12, and assuming that the U_h are differentiable (which we are assuming), the FOCS conditions for a maximum are both necessary and sufficient. So Pareto efficient points correspond to the solutions to

$$\alpha_h \pi_h(s) U'_h(x_{hs}) = \lambda_s \quad \text{and} \sum_s x_{hs} = W_s$$

without the nonnegativity constraint, and

$$\alpha_h \pi_h(s) U'_h(x_{hs}) \leq \lambda_s, \text{ with equality if } x_{hs} > 0, \quad \text{and} \sum_s x_{hs} = W_s$$

if we do impose the nonnegativity constraints.

Now consider how we might generate a solution for a given weighting vector (α_h). Take first the case where the nonnegativity constraints are not imposed. Pick a state s. For each nonnegative real number λ_s and for each h, look for solutions x to the equation

$$U'_h(x) = \lambda_s/(\alpha_h \pi_h(s)).$$

Let $X_{hs}(\lambda_s)$ be the set of solutions (which, clearly, depends on α_h and the parameter $\pi_h(s)$). (If $U_h'(-\infty) < \lambda_s/(\alpha_h \pi_h(s))$, let $X_{hs}(\lambda) = \{-\infty\}$. If $U_h'(\infty) > \lambda_s/(\alpha_h \pi_h(s))$, let $X_{hs}(\lambda) = \{\infty\}$.) Because U_h is concave and continuously differentiable, the sets X_{hs} are intervals that move continuously downwards in λ_s, meaning, they never overlap, and as λ increases, the set of solutions decreases in the sense that every member of the set for a smaller λ is greater than every member for a larger λ. Moreover, if U_h' is strictly concave, then $X_{hs}(\lambda)$ is always a singleton set and, therefore, describes a continuous, decreasing function. (If this isn't pretty close to obvious to you, you might benefit by going back to review the solutions of Problems 3.8 and 3.9. The mathematics are practically identical.) A solution to the FOCS conditions is obtained for state s when we find some λ, which will be λ_s, for which W_s is in the set-by-set sum (over h of the sets $X_{hs}(\lambda)$ and where each of these sets must contribute a finite element to the sum—you don't have an answer if $X_{hs}(\lambda) = \{-\infty\}$ for some individuals and $= \{\infty\}$ for others).

If the nonnegativity constraints are imposed, you follow the same process, except that in constructing the sets $X_{hs}(\lambda)$, you only look at nonnegative values for x, and if $U'(0) \leq \lambda/(\alpha_h \pi_h(s))$, then set $X_{hs}(\lambda) = \{0\}$.

The first bullet point concerns the case where every individual except perhaps for one is strictly risk averse. This means that each $X_{hs}(\lambda)$ will be singleton for every value of λ, except perhaps for one individual. Since the X_{hs} sets decline with increasing λ, there can be at most one λ that solves the FOCS conditions for each state (excluding the trivial case where the nonnegativity constraints are applied and $W_s = 0$ for some state), and for that one λ, which will be λ_s, the shares of all but at most one of the individuals are fixed. The share of the last individual is then fixed by the adding-up constraint (which must hold with equality).

The second bullet point is obvious from the FOCS conditions. Solutions depend on the weighting vector, the individual's subjective probability assessments, the shape of their utility functions, and W_s. The division of W_s into "shared ventures" and "private endowments" plays no role, and so this division can have no impact on efficient sharing rules.

Suppose that all the subjective probability assessments are the same: $\pi_h(s) = \pi_{h'}(s)$ for all h, h', and s. Then, given weights (α_h), the state-s maximization problem

$$\text{Maximize } \sum_h \alpha_h \pi_h(s) U_h(x_{hs}), \quad \text{subject to } \sum_h x_{hs} \leq W_s$$

and, possibly, nonnegativity constraints, is the same as

$$\text{Maximize } \sum_h \alpha_h U_h(x_{hs}), \quad \text{subject to the same constraints.}$$

Solutions depend in no fashion on the probability assessments, but only on the total amount of wealth, W_s, that can be distributed, the weighting vector (α_h), and the individual utility functions. (In terms of the FOCS conditions, the multipliers λ_s get scaled by the common probability assessment, but the set of solutions in terms of the sharing rule is unchanged by those probabilities.)

Suppose that all the subjective probability assessments are the same, and one party, say h_0, is risk neutral. Let U' be the constant slope of h_0's utility function. Then given a weighting vector (α_h), the value of the multiplier in state s is forced: $\lambda_s = \alpha_{h_0}\pi(s)U'$ (where I am writing $\pi(s)$ for the now common probability assessment). This is true independent of W_s—if there is a solution to the FOCS conditions, the solution must have this value of λ_s. (We are assuming that h_0 is not subject to the nonnegativity constraint, so her weighted-by-α-and-probability marginal utility must equal the mulitplier in any solution of the FOCS conditions.)

But then, for all the other (strictly risk averse) members of the syndicate, their net-of-endowment amounts are fixed by the weighting vector: x_{hs} must satisfy the equation

$$\alpha_h\pi(s)U_h'(x_{hs}) = \lambda_s = \alpha_{h_0}\pi(s)U' \quad \text{or} \quad U_h'(x_{hs}) = \frac{\alpha_{h_0}}{\alpha_h}U'.$$

(If individual h must satisfy the nonnegativity constraint, this equation is replaced by $x_{hs} = 0$ if $U_h'(0) \leq (\alpha_{h_0}/\alpha_h)U'$.) This solution is unique (strict concavity of the U_h) and is independent of W_s; the one risk-neutral individual has $X_{h_0s} = R$, so she is happy to soak up any residual risk after the rest of the syndicate gets their constant net-of-endowment shares. Note that the value of x_{hs} that h gets (net of endowment) varies in increasing fashion with h's weight relative to the weight on h_0; as α_{h_0}/α_h decreases, (as its reciprocal increases), the solution to $U_h'(x_{hs}) = (\alpha_{h_0}/\alpha_h)U'$ increases.

Finally, suppose $U_h(x) = -e^{-x/\tau_h}$ for each h. If nonnegativity constraints do not bind and there is a common probability assessment $\pi(s)$, the FOCS conditions are

$$\frac{\alpha_h}{\tau_h}e^{-x_{hs}/\tau_h} = \frac{\lambda_s}{\pi(s)} \quad \text{and} \quad \sum_s x_{hs} = W_s.$$

Write μ_s for $\lambda_s/\pi(s)$, and we can rewrite the first-order condition as

$$e^{-x_{hs}/\tau_h} = \frac{\tau_h\mu_s}{\alpha_h} \quad \text{or} \quad x_{hs} = -\tau_h \ln\left(\frac{\tau_h\mu_s}{\alpha_h}\right) = \tau_h \ln\left(\frac{\alpha_h}{\tau_h\mu_s}\right).$$

The adding-up constraint is then

$$\sum_h \tau_h \ln\left(\frac{\alpha_h}{\tau_h\mu_s}\right) = \sum_h \tau_h \ln\left(\frac{\alpha_h}{\tau_h}\right) - \ln(\mu_s)\sum_h \tau_h = W_s.$$

Let $K^* = \sum_h \tau_h \ln(\alpha_h/\tau_h)$, and let $T = \sum_h \tau_h$, and this becomes $K^* - T\ln(\mu_s) = W_s$, or $-\ln(\mu_s) = (W_s - K^*)/T$, where K^* is a constant (in s) that is parametrized by the vector of coefficients of risk tolerence (τ_h) and the weighting vector (α_h). And now we can go back and write

$$x_{hs} = \tau_h \ln\left(\frac{\alpha_h}{\tau_h}\right) - \tau_h \ln(\mu_s) = \tau_h \ln\left(\frac{\alpha_h}{\tau_h}\right) + \tau_h\left(\frac{W_s - K^*}{T}\right) =$$

$$\tau_h\left[\ln\left(\frac{\alpha_h}{\tau_h}\right) - \frac{K^*}{T}\right] \; + \; \frac{\tau_h}{T}W_s.$$

The term on the left-hand side of the plus sign is constant in s, depending (in a fairly complicated way) on the vectors (τ_h) and (α_h); moreover, these terms sum to zero.
And the term on the right-hand side of the plus sign, which says how h's net-of-endowment share varies with s, is as promised in the bullet point.