

EconS 594 – Ph.D. Industrial Organization

Fall 2026

Eco-labeling with Noisy Certification Tests: a Toy Model

Overview. This note works through a stripped-down version of Mason’s (2011) eco-labeling market. Each seller offers a single unit of a credence good, consumers value a green unit at 1 and a brown unit at 0, and the certifier runs a noisy test indexed by a single *precision* $\tau \in [0, 1]$. With these normalizations the analysis is short enough to fit on a whiteboard: **posterior beliefs equal equilibrium prices**, and the equilibrium conditions are linear in those beliefs. Seven steps follow: (a) the no-information benchmark, (b) the noisy test, (c) Bayes’ rule and prices, (d) application payoffs and the ranking $W_G > W_B$, (e) the equilibrium classes and the mixing probability in detail, (f) a numerical example, and (g) a note on what the toy model keeps and drops relative to Mason.

1. The Toy Model

Consumers and demand. Consumers value a green unit more than a brown one but cannot tell them apart directly. As in Mason (2011), demand is perfectly elastic. **Normalize** the two willingness-to-pay levels to

$$v_G = 1, \quad v_B = 0,$$

so the green “premium” is the whole willingness to pay. A consumer who believes a unit is green with probability μ pays the expected value $P(\mu) = \mu v_G + (1 - \mu) v_B = \mu$: *prices are beliefs*.

Sellers. There is a unit mass of sellers, each offering one unit. A fraction $\lambda \in (0, 1)$ are green and the complementary fraction $1 - \lambda$ are brown; a seller privately knows its own type. Because each seller brings a single unit to market, a seller’s payoff is simply the price it receives, net of any certification fee it pays. This single-unit structure is what keeps the equilibrium conditions linear in the beliefs.

Market structure. Within each observable segment (labeled or unlabeled) the market is competitive: the price equals consumers’ expected willingness to pay given the belief that a random unit in that segment is green. This is Mason’s rational-expectations pricing rule.

2. Step (a): The No-Information Benchmark

Suppose no eco-label is available. The whole market pools, and consumers pay the expected willingness to pay for a random unit,

$$P_0 = \lambda v_G + (1 - \lambda) v_B = \lambda. \quad (1)$$

Every seller, green or brown, receives $P_0 = \lambda$. This is the reference price against which the label is measured, and the outcome the market reverts to whenever the fee is too high for anyone to seek certification.

3. Step (b): The Noisy Certification Test

A certifying agency offers a test. A seller that applies pays a fee $A \geq 0$ and receives a public verdict (pass or fail). We index the test by a single *precision* $\tau \in [0, 1]$: a green seller passes with probability ϕ_G and a brown seller with probability ϕ_B , where

$$\phi_G \equiv \Pr(\text{pass} \mid G) = \frac{1 + \tau}{2}, \quad \phi_B \equiv \Pr(\text{pass} \mid B) = \frac{1 - \tau}{2}. \quad (2)$$

At $\tau = 0$ the test is uninformative ($\phi_G = \phi_B = \frac{1}{2}$); at $\tau = 1$ it is perfect ($\phi_G = 1, \phi_B = 0$). A green seller fails with probability $1 - \phi_G = \frac{1 - \tau}{2}$ (a *type I* error, or false rejection), and a brown seller passes with probability $\phi_B = \frac{1 - \tau}{2}$ (a *type II* error, or false acceptance). The two error rates are equal, so τ is the single dial that sets test quality, with informativeness $\phi_G - \phi_B = \tau$. Table 1 collects the four outcomes conditional on applying, and Table 2 shows its two extremes. Figure 1 plots the two pass probabilities against τ , and Figure 2 shows the same four flows as a tree.

	Labeled (pass)	Unlabeled (fail)
Green seller	$\frac{1 + \tau}{2}$	$\frac{1 - \tau}{2}$ (type I error)
Brown seller	$\frac{1 - \tau}{2}$ (type II error)	$\frac{1 + \tau}{2}$

Table 1: Certification outcomes by seller type, conditional on applying. Under the symmetric precision test (2), the diagonal (correct classification) carries mass $\frac{1 + \tau}{2}$ and each off-diagonal error carries mass $\frac{1 - \tau}{2}$; a more precise test (higher τ) moves mass onto the diagonal.

$\tau = 0$	Labeled	Unlabeled
Green seller	$\frac{1}{2}$	$\frac{1}{2}$
Brown seller	$\frac{1}{2}$	$\frac{1}{2}$

(a) Uninformative test.

$\tau = 1$	Labeled	Unlabeled
Green seller	1	0
Brown seller	0	1

(b) Perfect test.

Table 2: The two extremes of the confusion matrix. At $\tau = 0$ (panel a) every cell equals $\frac{1}{2}$: the label is pure noise, so the two rows are identical and a pass says nothing about type. At $\tau = 1$ (panel b) the diagonal is 1 and both errors vanish: the label separates green from brown perfectly.




Figure 1: The two pass probabilities as functions of precision τ : a green seller passes with probability $\phi_G = \frac{1 + \tau}{2}$ (rising, green) and a brown seller with probability $\phi_B = \frac{1 - \tau}{2}$ (falling, red). They coincide at $\frac{1}{2}$ when $\tau = 0$ (uninformative) and separate to 1 and 0 when $\tau = 1$ (perfect). The vertical gap between them is the test's informativeness, $\phi_G - \phi_B = \tau$.




Figure 2: The four flows through the noisy test. Green sellers pass with the higher probability ϕ_G ; brown sellers pass with the lower probability ϕ_B . Failures and non-applicants are pooled together in the *unlabeled* segment.

4. Step (c): Posterior Beliefs and Prices

Suppose (for now) that every seller applies. Because each seller brings one unit, the mass of green units in a segment relative to the total is a plain Bayes ratio, and the two posteriors are

$$\mu \equiv \Pr(G \mid \text{labeled}) = \frac{\lambda\phi_G}{\lambda\phi_G + (1-\lambda)\phi_B}, \quad (3)$$

$$\nu \equiv \Pr(G \mid \text{unlabeled}) = \frac{\lambda(1-\phi_G)}{\lambda(1-\phi_G) + (1-\lambda)(1-\phi_B)}. \quad (4)$$

Substituting the precision test (2),

$$\mu = \frac{\lambda(1+\tau)}{\lambda(1+\tau) + (1-\lambda)(1-\tau)}, \quad \nu = \frac{\lambda(1-\tau)}{\lambda(1-\tau) + (1-\lambda)(1+\tau)}. \quad (5)$$

At $\lambda = \frac{1}{2}$ these collapse to $\mu = \frac{1+\tau}{2}$ and $\nu = \frac{1-\tau}{2}$, the pass probabilities themselves.

Prices equal posteriors. Within each segment competition sets the price to the expected willingness to pay, which under $v_G = 1$ and $v_B = 0$ is exactly the posterior:

$$P_c = \mu, \quad P_u = \nu, \quad P_0 = \lambda. \quad (6)$$

Sorting. Any precision $\tau > 0$ makes $\phi_G > \phi_B$, which translates into $\mu > \lambda > \nu$: an informative test raises the green share (and hence the price) in the labeled segment above the population share and lowers it in the unlabeled segment. Using $\phi_G - \phi_B = \tau$,

$$\mu - \lambda = \frac{\lambda\phi_G - \lambda[\lambda\phi_G + (1-\lambda)\phi_B]}{\lambda\phi_G + (1-\lambda)\phi_B} = \frac{\lambda(1-\lambda)\tau}{\lambda\phi_G + (1-\lambda)\phi_B} > 0,$$

and $\lambda - \nu > 0$ by an analogous computation.

Result 1 (sorting). The label puts the two segments on opposite sides of the prior, $\nu < \lambda < \mu$, with a gap proportional to the precision τ and to the mixing factor $\lambda(1-\lambda)$.

5. Step (d): Application Payoffs and Ordered Incentives

A seller of type $k = G, B$ that applies is labeled with probability ϕ_k , earning the certified price $P_c = \mu$, and unlabeled with probability $1 - \phi_k$, earning $P_u = \nu$; it pays the fee A either

way. If it does not apply it is unlabeled for sure and earns ν . The gain from applying is therefore

$$W_k \equiv [\phi_k \mu + (1 - \phi_k) \nu - A] - \nu = \phi_k (\mu - \nu) - A, \quad (7)$$

which is *linear* in the price gap $\mu - \nu$. Subtracting the two gains, and using $\phi_G - \phi_B = \tau$,

$$W_G - W_B = (\phi_G - \phi_B)(\mu - \nu) = \tau (\mu - \nu) > 0. \quad (8)$$

Ordered incentives. Green sellers have a strictly stronger incentive to seek certification than brown sellers, because their higher pass probability multiplies the same price gap into a larger gain. As a corollary, no equilibrium can have some brown sellers applying while no green sellers do; if brown sellers are in, greens are too.

6. Step (e): The Equilibrium Classes

Depending on the size of the fee A , the equilibrium falls into one of four regions, corresponding to three kinds of Perfect Bayesian equilibrium (PBE): a *separating* PBE, a *partially separating* PBE, and a *pooling* PBE that appears in two forms (both sellers apply, or no seller applies). The cutoffs come out in closed form.

Pooling PBE (both sellers apply). Both types apply, so the market cannot tell green from brown among applicants. This requires $W_B \geq 0$, i.e.,

$$A \leq \phi_B (\mu - \nu) \equiv \bar{A}_{\text{pool}}, \quad (9)$$

with μ, ν the pooling posteriors (3)–(4).

Separating PBE (green apply, brown do not). Since only greens apply, every labeled unit is green, so $\hat{\mu} = 1$ and $\hat{P}_c = 1$. The unlabeled posterior contains greens that failed plus all browns:

$$\hat{\nu} = \frac{\lambda(1 - \phi_G)}{\lambda(1 - \phi_G) + (1 - \lambda)}, \quad \hat{P}_u = \hat{\nu}. \quad (10)$$

Because the certified price is now 1, the price gap is $1 - \hat{\nu}$, and the separating range on the fee is

$$\underbrace{\phi_B (1 - \hat{\nu})}_{\text{brown IC: do not apply}} \leq A \leq \underbrace{\phi_G (1 - \hat{\nu})}_{\text{green IC: do apply}}. \quad (11)$$

Since $\phi_G > \phi_B$, this interval has positive length.

Partially separating PBE (green apply, brown mix). For fees in the gap

$$\bar{A}_{\text{pool}} < A < \phi_B (1 - \hat{\nu}), \quad (12)$$

brown sellers mix so as to be indifferent between applying and not applying. The mixing probability $\lambda_B \in (0, 1)$ solves the indifference condition

$$\phi_B (\mu^* - \nu^*) = A, \quad (13)$$

where the labeled and unlabeled posteriors, when the green seller applies for sure and the brown seller applies with probability λ_B , are

$$\mu^* = \frac{\lambda \phi_G}{\lambda \phi_G + (1 - \lambda) \lambda_B \phi_B}, \quad (14)$$

$$\nu^* = \frac{\lambda(1 - \phi_G)}{\lambda(1 - \phi_G) + (1 - \lambda)(1 - \lambda_B \phi_B)}. \quad (15)$$

These reduce to the pooling posteriors (3)–(4) when $\lambda_B = 1$ and to the separating posteriors ($\mu^* = 1$, $\nu^* = \hat{\nu}$) when $\lambda_B = 0$. This regime deserves a closer look, since the mixing probability is what makes the label only partly trustworthy.

The mixing probability lies strictly between 0 and 1. Write the brown seller's gain from applying as a function of the mixing probability,

$$g(\lambda_B) \equiv \phi_B (\mu^* - \nu^*), \quad (16)$$

so the indifference condition (13) is $g(\lambda_B) = A$. As more brown sellers apply, the labeled pool is diluted and the unlabeled pool is cleansed: in (14) the numerator of μ^* is fixed while its denominator grows, so μ^* *falls* in λ_B ; in (15) the term $(1 - \lambda_B \phi_B)$ shrinks, so ν^* *rises*. Both moves shrink the gap $\mu^* - \nu^*$, so g is strictly decreasing, from

$$g(0) = \phi_B (1 - \hat{\nu}) = \underline{A}_{\text{sep}} \quad \text{down to} \quad g(1) = \phi_B (\mu - \nu) = \bar{A}_{\text{pool}}.$$

A partially separating fee lies strictly between these two cutoffs, $\bar{A}_{\text{pool}} < A < \underline{A}_{\text{sep}}$, so by continuity and strict monotonicity the equation $g(\lambda_B) = A$ has a *unique* solution $\lambda_B \in (0, 1)$. Outside the gap there is no interior solution: for $A \leq \bar{A}_{\text{pool}}$ the brown seller strictly prefers to apply ($\lambda_B = 1$, pooling), and for $A \geq \underline{A}_{\text{sep}}$ it strictly prefers not to ($\lambda_B = 0$, separation).

Comparative static in the fee. Differentiating $g(\lambda_B) = A$ and using $g' < 0$,

$$\frac{\partial \lambda_B}{\partial A} = \frac{1}{g'(\lambda_B)} < 0, \quad (17)$$

so a higher fee lowers the fraction of brown sellers that apply. The intuition is a dilution feedback: a higher fee makes the label less attractive to a brown seller, so fewer apply; but fewer brown applicants make the label *cleaner* (a higher μ^* and a wider price gap), which pushes the gain from applying back up until the marginal brown seller is again indifferent. At $\lambda = \frac{1}{2}$ and $\tau = \frac{4}{5}$, where (14)–(15) simplify to $\mu^* = \frac{9}{9 + \lambda_B}$ and $\nu^* = \frac{1}{11 - \lambda_B}$, the solution runs monotonically from $\lambda_B = 1$ at $A = \frac{2}{25}$ to $\lambda_B = 0$ at $A = \frac{1}{11}$, as Table 3 reports.

Fee A	Brown application probability λ_B
$\frac{2}{25} = 0.080$	1 (pooling boundary)
0.082	0.80
0.084	0.61
0.086	0.43
0.088	0.25
$\frac{1}{11} \approx 0.091$	0 (separating boundary)

Table 3: The brown seller’s equilibrium application probability λ_B across the partially separating gap, at $\lambda = \frac{1}{2}$ and $\tau = \frac{4}{5}$. As the fee rises from $\frac{2}{25}$ to $\frac{1}{11}$, λ_B falls from 1 to 0.

Comparative static in precision. The effect of τ on λ_B is not monotone, because a more precise test pulls in two directions. It lowers the brown pass probability $\phi_B = \frac{1 - \tau}{2}$, which discourages brown application, but it widens the price gap $\mu^* - \nu^*$, which encourages it. Either force can dominate, so λ_B can rise or fall with τ , the same non-monotonicity Mason (2011) reports. The gap itself, $(\bar{A}_{\text{pool}}(\tau), \underline{A}_{\text{sep}}(\tau))$, also shifts and widens with τ , so a fixed fee can move out of the partially separating region altogether as the test becomes more precise.

What this means for the green seller. The green seller applies for sure. At the brown seller’s indifference, $\mu^* - \nu^* = \frac{A}{\phi_B}$, so the green seller’s net gain from applying is

$$\phi_G(\mu^* - \nu^*) - A = A \frac{\phi_G - \phi_B}{\phi_B} = \frac{2A\tau}{1 - \tau} > 0, \quad (18)$$

strictly positive and increasing in both the fee and the precision. The brown seller’s mixing

is exactly what pins the label's value: it leaves the brown seller indifferent while the green seller, passing the test more often, still strictly profits. A more precise test widens this wedge, so the green seller's advantage from certification grows with τ even though the brown seller is held to zero rents.

Pooling PBE (no seller applies). When the fee clears the upper separating cutoff, $A > \phi_G(1 - \hat{\nu})$, not even a green seller applies. Both types pool on not applying, and the market reverts to the no-information benchmark of Step (a).

The map from test fee to equilibrium. As A rises from zero the equilibrium moves through the pooling PBE in which both sellers apply, the partially separating PBE, the separating PBE, and finally the pooling PBE in which no seller applies. The cutoffs (9), (11), and (12) are ordered strictly, because $\phi_G > \phi_B$ and the pooling posterior μ is strictly less than the separating posterior $\hat{\mu} = 1$.

The equilibrium map. At $\lambda = \frac{1}{2}$ the pooling posteriors reduce to $\mu = \frac{1 + \tau}{2}$ and $\nu = \frac{1 - \tau}{2}$, so $\mu - \nu = \tau$ and $\hat{\nu} = \frac{1 - \tau}{3 - \tau}$. The three cutoffs of (9), (11), and (12) then become

$$\bar{A}_{\text{pool}}(\tau) = \frac{\tau(1 - \tau)}{2}, \quad \underline{A}_{\text{sep}}(\tau) = \frac{1 - \tau}{3 - \tau}, \quad \bar{A}_{\text{sep}}(\tau) = \frac{1 + \tau}{3 - \tau}. \quad (19)$$

These are the same cutoffs that Mason (2011) obtains. Plotting the three curves in the (τ, A) plane carves it into the four regions of Figure 3: the shaded separating region between $\underline{A}_{\text{sep}}$ and $\bar{A}_{\text{sep}}$, the pooling region (both apply) below $\bar{A}_{\text{pool}}$, the partially separating wedge between them, and the pooling region (no seller applies) above $\bar{A}_{\text{sep}}$. The vertical slice at $\tau = \frac{4}{5}$ gives the numerical example of Step (f): $\bar{A}_{\text{pool}} = \frac{2}{25}$, $\underline{A}_{\text{sep}} = \frac{1}{11}$, and $\bar{A}_{\text{sep}} = \frac{9}{11}$.

7. Step (f): A Numerical Example

Take

$$\lambda = \frac{1}{2}, \quad \tau = \frac{4}{5}, \quad A = \frac{1}{20},$$

so from the precision test (2) the pass probabilities are $\phi_G = \frac{9}{10}$ and $\phi_B = \frac{1}{10}$ (a type I and a type II error rate of $\frac{1}{10}$ each).

Prices. The no-label price is $P_0 = \lambda = \frac{1}{2}$. At $\lambda = \frac{1}{2}$ the posteriors equal the pass




Figure 3: The equilibrium map in the (τ, A) plane, at $\lambda = \frac{1}{2}$. The three cutoffs of (19) split the plane into the four regions of Step (e): a separating PBE (only green apply, shaded), a partially separating PBE (brown randomizes), and the pooling PBE in its two forms (both sellers apply at low fees, no seller applies at high fees). The separating region widens with precision; the both-apply pooling region is a thin sliver that is largest at intermediate precision. The dotted vertical slice at $\tau = \frac{4}{5}$ is the numerical example of Step (f).

probabilities, so

$$\mu = \frac{1 + \tau}{2} = \frac{9}{10}, \quad \nu = \frac{1 - \tau}{2} = \frac{1}{10}, \quad \mu - \nu = \frac{4}{5}.$$

Prices are the posteriors: $P_c = \frac{9}{10}$ and $P_u = \frac{1}{10}$, straddling $P_0 = \frac{1}{2}$.

Which equilibrium? The pooling cutoff (9) is

$$\bar{A}_{\text{pool}} = \phi_B(\mu - \nu) = \frac{1}{10} \cdot \frac{4}{5} = \frac{2}{25} = 0.08.$$

Since $A = \frac{1}{20} = 0.05 < 0.08 = \bar{A}_{\text{pool}}$, we are in the pooling PBE: both types apply.

Application gains. From (7), $W_k = \phi_k(\mu - \nu) - A = \phi_k \cdot \frac{4}{5} - \frac{1}{20}$:

$$W_G = \frac{9}{10} \cdot \frac{4}{5} - \frac{1}{20} = \frac{36}{50} - \frac{1}{20} = \frac{72 - 5}{100} = \frac{67}{100}, \quad W_B = \frac{1}{10} \cdot \frac{4}{5} - \frac{1}{20} = \frac{8 - 5}{100} = \frac{3}{100}.$$

Both are strictly positive, and $W_G > W_B$, as required. Raising the fee above $\frac{2}{25}$ would push

the market into the partially separating region of Table 3, and above $\frac{1}{11}$ into a separating PBE.

8. Step (g): What the Toy Model Keeps and Drops

Every one of Mason’s positive conclusions about the label survives the normalizations:

1. **Bayesian sorting**, $\nu < \lambda < \mu$, with a gap proportional to the precision τ (Result 1).
2. **Ordered incentives**, $W_G > W_B$, so brown sellers never apply unless green sellers do (Step d).
3. **Four equilibrium regions** indexed by the fee: a pooling PBE in which both apply (low A), a partially separating PBE (intermediate A), a separating PBE (high A), and a pooling PBE in which no seller applies (prohibitive A), with the same cutoffs $\underline{A}_{\text{sep}} = \frac{1 - \tau}{3 - \tau}$ and $\bar{A}_{\text{sep}} = \frac{1 + \tau}{3 - \tau}$ that Mason obtains, shown in Figure 3.
4. **The mixing probability is interior and well-behaved**: it is unique, decreasing in the fee (Table 3), and non-monotone in precision, and it leaves the green seller with a strictly positive gain $\frac{2A\tau}{1 - \tau}$ that grows with τ (Step e).

What the toy model drops, relative to Mason (2011) and to the certifier-design literature, is threefold: (i) sellers here bring a single unit, so there is no choice of output and no producer-surplus welfare accounting; (ii) the two error rates are tied together by the symmetric test (2) at $\frac{1 - \tau}{2}$, whereas Mason lets the two pass rates move independently; and (iii) the precision τ is fixed, rather than chosen by a certifier or a regulator who trades the value of information against the cost of a more accurate test. Each can be reinstated one at a time; the payoff of the toy model is that the sorting mechanism and the four equilibrium regions come through with almost no algebra.

Reference

Mason, C. F. (2011). “Eco-Labeling and Market Equilibria with Noisy Certification Tests.” *Environmental and Resource Economics*, 48(4), pp. 537–560. <https://doi.org/10.1007/s10640-010-9402-5>