

Supplementary Mathematical Appendix

Human-in-the-Loop Tax Auditing

Kun Huang
Wuhan University

October 3, 2026

This supplement provides the proofs and restrictions of the eight analytical results in the accompanying manuscript. Counterexamples are mathematical constructions, not taxpayer observations. The computational supplement supplies separate numerical checks; simulations do not prove the results.

Common environment. State $\theta \in \{0, 1\}$ has calibrated posterior p given retained AI information. Automated processing has value $pg_A - (1 - p)f - \bar{c}_A$, and human investigation supplies conditional signal probabilities α, β before choosing release or intensive processing with true-state payoff g_H . The welfare primitives are $g_a = \mu r_a R - b_1$, $f = b_0$, $\bar{c}_A = (1 + \mu)c_A$, $k_H = (1 + \mu)c_H + b_H$; fiscal primitives are $g_a = r_a R$, $f = 0$, $\bar{c}_A = c_A$, $k_H = c_H$. Actual administrative costs and real burdens are nonnegative. Unit-slot allocation has fixed separable gains. Conditional error coupling is state symmetric. The deterrence extension is a separate commitment model and explicitly restricts parameters to interior concealment choices.

1 Mathematical appendix

1.1 Allocation, existence, and capacity value

For any selected set S , optimal nonhuman decisions contribute $\sum_i B_{Ai}$. Review adds $\sum_{i \in S} v_i$ and consumes $|S|$ slots. Removing a negative increment strictly improves the objective. If an unselected case has a larger gain than a selected case, their exchange preserves capacity and strictly improves value. Repeated exchanges yield the largest positive gains, proving Theorem 1. Equal gains permit multiple deterministic optima. Neither perfect information nor universally superior human performance is needed.

For the fractional relaxation, the dual objective is

$$\min_{\lambda \geq 0} \left\{ \sum_i B_{Ai} + \lambda H + \sum_i (v_i - \lambda)_+ \right\}.$$

At a binding integer cutoff without ties, any λ between the next unselected positive gain and the smallest selected gain supports the allocation. At slack capacity the multiplier is zero. The reported next-slot value corresponds to the lower endpoint of this interval. The continuous atomless problem has the analogous dual with an integral. Integrable gains ensure finite value; choosing all gains above a quantile threshold and randomizing on a positive-mass tie implements any binding capacity. The primal and dual values agree by the explicit threshold construction.

The value with H integer slots is the baseline plus the sum of the first H nonnegative order statistics. Its forward difference is the next order statistic, proving Proposition 2. A fixed tie convention gives a nested sequence. Fractional interpolation is concave because its slopes are ordered. A deterministic value function evaluated at the floor of a real-valued capacity is a step function; the concavity claim concerns integer differences or the fractional relaxation. If case duration differs, neither the exchange proof nor raw-gain sorting suffices. For example, a value-10 case requiring two slots loses to two value-6 cases requiring one each, despite having the largest individual gain.

1.2 The decomposition and homogeneous review interval

By posterior consistency and iterated expectations, the gross reviewed payoff is $\mathbb{E}B_H(p_Z)$. Adding and subtracting $B_H(p)$ gives equation (4). Convexity of the maximum of affine payoffs implies $\mathbb{E}B_H(p_Z) \geq B_H(p)$; equivalently, the authority can ignore the signal. Because $g_H \geq g_A$ in the maintained technology ordering, $B_H(p) \geq B_A(p)$. These nonnegative components do not imply positive net review value once k_H is charged.

For identical technologies, write the terminal processing value as $pG - (1 - p)L$. The positive and negative human signals switch that decision at

$$p_- = \frac{\beta L}{\alpha G + \beta L}, \quad t = \frac{L}{G + L}, \quad p_+ = \frac{(1 - \beta)L}{(1 - \alpha)G + (1 - \beta)L}.$$

When $G, L > 0$ and $\alpha > \beta$, one has $p_- < t < p_+$, with limiting weak inequalities at boundary channels. Below p_- both signals lead to release, and above p_+ both lead to processing. Gross information value is therefore

$$I(p) = \begin{cases} 0, & p \leq p_-, \\ \alpha G p - \beta L(1 - p), & p_- < p \leq t, \\ (1 - \beta)L(1 - p) - (1 - \alpha)G p, & t < p < p_+, \\ 0, & p \geq p_+. \end{cases}$$

The middle expressions have positive and negative slopes, respectively, and agree at t , where the value is $(\alpha - \beta)GL/(G + L)$. This proves Proposition 3. If the peak exceeds $k_H + \lambda$, the strict review interval has endpoints

$$\ell = \frac{\beta L + k_H + \lambda}{\alpha G + \beta L}, \quad u = \frac{(1 - \beta)L - k_H - \lambda}{(1 - \alpha)G + (1 - \beta)L}.$$

Below ℓ the nonhuman optimum releases; above u it processes. At equalities either optimal branch can be used consistently with capacity. If $G \leq 0$ and $L \geq 0$, processing is never valuable and this nondegenerate interval characterization does not apply.

For the disconnected example in Results, human gross value is piecewise linear with breakpoints $1/19$ and $9/11$, while the direct-processing baseline changes at $1/2$. On the first active human segment, $v = 1.9p - 0.93$ for $p < 1/2$ and $v = 0.07 - 0.1p$ for $1/2 < p < 9/11$. Above $9/11$, $v = p - 0.83$. Setting each expression greater than zero yields $(0.93/1.9, 0.7) \cup (0.83, 1]$. The exact lower endpoint is $93/190$, and the reported decimal is rounded. This verifies that adding a recovery advantage can produce $N-H-A-H$ even with homogeneous costs and signal probabilities.

1.3 Refinement, costs, and selected sets

If a new human signal can reproduce the old one through state-independent garbling conditional on X , the authority can apply that garbling and follow its old terminal decision. The new signal

therefore cannot lower optimal gross value, proving the refinement part of Proposition 4. This is a statement about a specified information order; an increase in an aggregate accuracy statistic alone need not imply it. A common added review cost d changes each gain to $v_i - d$. Pairwise differences are unchanged. With the same tie ordering, selected cases either remain or exit as positive demand contracts. Heterogeneous improvements can instead reorder cases.

For example, two cases with $p = 1/2$, $G = L = 1$, and review cost 0.1 have symmetric human qualities 0.70 and 0.65. Net gains are 0.10 and 0.05. With one slot, the first is reviewed. Improving the second quality to 0.80 gives gain 0.20 and displaces the first, although no individual gain fell. Binding capacity prevents a strict increase in review count, and selected sets are not nested under this heterogeneous refinement.

1.4 Error dependence and the clone restriction

The joint error table in Methods follows from the two marginal errors and their joint intersection c . Conditioning on B gives the stated α, β . The four-cell probabilities must be imposed identically conditional on each state; overall marginal accuracies would not otherwise identify the conditional channel. When $q = h$, the conditional channel at dependence ρ is a mixture: with probability $1 - \rho$ use independent human evidence, and with probability ρ copy the already observed B .

For $\rho_2 > \rho_1$, force the former human output to equal B with probability $(\rho_2 - \rho_1)/(1 - \rho_1)$. This garbling is independent of the state conditional on B and produces precisely the channel at ρ_2 . The previous refinement argument proves monotone information loss. At $\rho = 1$ there is no posterior change, hence $I = 0$ and $v = P - k_H$. The procedure can still add value, proving Proposition 5 in its qualified form.

For the unequal-quality counterexample, equation (5) yields $\alpha = 0.7 + (7/90)\rho$ and $\beta = 0.3 + 0.7\rho$. At $p = 1/2$ and $G = L = 1$, the two branch contributions are opposite, so the sum of their positive parts is $|\alpha - \beta|/2$. Substitution yields the expression in Results. At dependence exceeding $9/14$, $\alpha < \beta$, and acting on the conventionally negative signal is optimal. Both marginal accuracies are unchanged. This lawful counterexample refutes the unrestricted correlation comparative static rather than contradicting the equal-quality theorem.

1.5 AI information, fiscal disagreement, and limiting cases

For Proposition 6, take any optimal policy using the former AI information. Under richer information, ignore the added evidence and use the retained former information, select the same review cases, and make the same terminal decisions. Capacity, physical human evidence, and payoffs are unchanged. The old expected value is feasible, so the new optimum weakly exceeds it. This ex ante argument relies on literal retention and the unchanged joint signal environment; it is not a pointwise posterior comparison. Marginal Blackwell dominance of the machine signal alone, with only the state-human marginal law fixed, would not suffice because dependence with human evidence could change.

To show substitution, take prior $1/2$, $G = L = 1$, perfect human evidence, and review cost 0.1. An uninformative machine gives review gain 0.4. A perfect machine gives gain -0.1 , and positive review demand disappears. To show complementarity, take prior 0.95, human symmetric quality 0.8, and the same payoffs and review cost. An uninformative machine gives zero gross information value, because even unfavorable human evidence cannot change the action. A machine with symmetric quality 0.9 generates unfavorable outputs with probability 0.14 and posterior $19/28$. Conditional human gross value there is $17/140$, giving net gain $3/140$. At a capacity below 0.14, the slot value rises from zero to $3/140$. Thus total information improvement does not determine the sign of review demand or shadow value.

For Proposition 7, if $\mu = 0$, then $g_A = g_H = -b_1 \leq 0$, $f = b_0 \geq 0$, and all processing values are nonpositive. The human branch optimally closes after every signal and pays k_H . Release dominates costly intervention. To construct disagreement between objectives, take two cases with perfect human evidence, actual $c_A = 0.1$, $c_H = 0.5$, and $b_0 = b_H = 0$. Case 1 has $(p, R, r_A, r_H, b_1) = (0.9, 100, 0.5, 1, 10)$; case 2 has $(0.5, 40, 0.1, 1, 0)$. Fiscal gains are 44.51 and 17.55, selecting case 1. At $\mu = 0.1$, financed welfare costs are 0.11 and 0.55. Welfare gains are -0.55 and 1.305 , selecting case 2. State probabilities and human accuracy are unchanged.

1.6 Deterrence response and optimal commitment

The private objective is strictly concave in e when $\kappa > 0$, and its interior first-order condition gives $e = (T - FD)/\kappa$. The imposed restrictions $0 < T - F(d_A + \delta)$ and $T < \kappa$ ensure $0 < e < 1$ for every $t \in [0, 1]$. Also require $0 \leq d_A \leq d_A + \delta \leq 1$, $F \geq T > 0$, and $\delta \geq 0$. The total private sanction includes recovered tax, while additional penalty receipts have no separate premium in the specified welfare objective.

Because $D' = \delta$ and $e' = -F\delta/\kappa$, differentiating equation (6) gives

$$W' = Fe\delta + \mu T\{-e' + e'D + e\delta\} - \Delta c = \delta[Fe + \mu T\{e + F(1 - D)/\kappa\}] - \Delta c.$$

Differentiating again yields $W'' = -\delta^2(F^2 + 2\mu TF)/\kappa$. Concavity follows for $\mu \geq 0$. The Lagrangian separates by type and gives $M_j(t_j) = \lambda$ for interior choices, $M_j(0) \leq \lambda$ for unused types, and $M_j(1) \geq \lambda$ for fully reviewed types. With $a_j = M_j(0)$ and $s_j = \delta_j^2(F_j^2 + 2\mu T_j F_j)/\kappa_j$, the solution for $s_j > 0$ is

$$t_j(\lambda) = \min\{1, \max\{0, (a_j - \lambda)/s_j\}\}.$$

If the resulting capacity at $\lambda = 0$ is feasible, it is the optimum. Otherwise a scalar root enforces $\sum_j m_j t_j = K$. Zero detection increments and nonnegative review costs optimally receive no review. This establishes Proposition 8 without relying on simulation.

For the two-type example, the marginal schedules are $M_A(t) = 0.125 - 0.0075t$ and $M_B(t) = 0.260 - 0.060t$. With masses one-half each and $K = 0.25$, setting $t_B = 0.5$ exhausts capacity and gives multiplier 0.23. Since $M_A(0) < 0.23$, type A remains unreviewed. The frozen-response recovery ranking is reversed. Both concealment choices remain interior throughout the feasible interval, so no omitted clipping or discontinuity drives the example.