



Introduction & Motivation

Kenyan Judiciary introduced Court Annexed Mediation (CAM) where:

- Civil cases assigned to court appointed mediators
- If mediation fails to reach a resolution, cases go to trial in front of a judge
- CAM enables faster case resolution, reduces administrative costs, and provides a less adversarial dispute-resolution process.

Current Mechanism:


- CAM officers assign mediators case-by-case at own discretion.
- Mediators licensed by court station and case type.
- Rules: accreditation + max 3 open cases per mediator

Observations and Shortcomings:

- 58% of cases return unresolved
- Success rates vary widely across mediators
- Mediator capability judged subjectively by CAM officers
- Capacity constraints often violated by substantial numbers

Mediator Capability Modeling:

- We use Value Added Model [1]
- Reveals substantial heterogeneity in mediator skills
- 1 SD increase in VA $\Rightarrow$ 11 pp increase in probability of agreement




Objective

Our goal is to propose an automated case allocation system that:

- Maximizes case agreement rate while reducing mediator overload

Key questions:

- How much value do mediators add?**
 - Historical data suggests agreement rate varies by mediators. Accurately assessing their contribution is crucial for making allocation choices.
 - Mediator experiences can vary and new mediators can enter with no history
 - Value added by mediators needs to be learned online
 - As such, the system must trade off **exploration and exploitation** to find good mediators and utilize them
- Is there a benefit to holding back mediators occasionally?**
 - Capacity limits make mediator assignment a long-term tradeoff
 - Greedy assignment isn't always optimal under capacity limits
 - High opportunity cost assignment: Mediators who can take cases in multiple (type-stations) cells should be held back in mediator-rich cells, allowing for better capacity management in mediator scarce cells.




- How to allow inevitable constraint violations efficiently:**
 - High case-volume may make capacity constraints impossible to maintain
 - A trade off exists between constraint violation and case resolution rate

We propose a novel allocation algorithm that deals with all 3 issues

Experiments

Stylized Scenarios



Algorithm	λ	Known VA (SMaRT(GS))	Learning + Mean VA (SMaRT(E))	Learning + Sampled VA (SMaRT(S))	Algorithm	λ	Known VA (SMaRT(GS))	Learning + Mean VA (SMaRT(E))	Learning + Sampled VA (SMaRT(S))
		Agreement Rate	OCMD	Agreement Rate			Agreement Rate	OCMD	Agreement Rate
Least Load	—	0.406	25.13	—	Least Load	—	0.380	4.13	—
Greedy*	—	0.439	41.71	0.425	Greedy*	—	0.406	5.18	0.394
TS*	—	—	—	36.31	TS*	—	—	5.06	0.392
				0.404					4.82
SMaRT	0.01	0.453	94.12	0.438	SMaRT	0.01	0.452	94.30	0.417
	0.05	0.449	54.34	0.430		0.05	0.440	43.90	0.403
	0.10	0.444	44.85	0.427		0.10	0.430	22.91	0.398
	0.50	0.439	38.53	0.422		0.50	0.410	5.43	0.394
				33.04					5.54
Upper Bound	—	0.460	132.15	—	Upper Bound	—	0.460	132.15	—

Table 1: Scenario 1 Results

Table 2: Scenario 2 Results




Figure 1: An example of case allocations (with known med VAs) in Scenario 1 by SMaRT

- Tradeoff: Agreement Rate $\Leftrightarrow$ Overload
- Low λ : High Agreement Rate + More Overload for high VA mediators
- High λ : Cases distributed more equitably + Agreement Rate is lower
- Tuning λ : More 'Greedy' like, 'Best Mediator' like; or a hybrid in between
- Simulations are run 128 times for each setting, and the results are averaged




Figure 2: Mean shadow price by mediator for varying λ

References

- [1] Raj Chetty, John N. Friedman, and Jonah E. Rockoff. "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates." American Economic Review, 2014.
- [2] Aleksandrs Slivkins. "Introduction to Multi-Armed Bandits." Foundations and Trends in Machine Learning, 2019.
- [3] Ashwinkumar Badanidiyuru, Robert Kleinberg, and Aleksandrs Slivkins. "Bandits with Knapsacks." Journal of the ACM, 2018.
- [4] Martino Bernasconi, Matteo Castiglioni, Andrea Celli, and Federico Fusco. "Bandits with Replenishable Knapsacks: The Best of Both Worlds." ICLR, 2024.
- [5] Shipra Agrawal and Nikhil R. Devanur. "Bandits with Concave Rewards and Convex Knapsacks." EC, 2014.
- [6] Aditya Mate, Jackson A. Killian, Haifeng Xu, Andrew Perrault, and Milind Tambe. "Collapsing Bandits and Their Application to Public Health Interventions." NeurIPS, 2020.

Related Work

MABs [2] have been used extensively in social resource allocation settings [6, 11]

Existing methods in the literature that are ill suited for this problem:

- Classic MABs can learn mediator quality but ignore capacity and accreditation constraints [2].
- BwK [3], it's replenishable extension [4], and generalization with penalties [5] cannot accommodate immediate assignment and soft constraint violations.

Mathematical programming has often been used for bipartite resource matching:

- [7-9] use Linear Programming (LP) to model online bipartite matching
- OM-RR-PA [10] uses LP and supports reusable resources, but assumes delayed assignment, hard capacity, and known service times.

Selecting Mediators that are Right for the Task (SMaRT)

Allocation: With i) full knowledge of future cases,

ii) Knowledge of mediator true Value Addeds (VA), and

iii) Relaxation allowing fractional allocations,

The allocation problem can be expressed with a QP.

To overcome imperfect knowledge of future arrivals:

- Shadow cases (V_{sh}) are sampled from case-arrival distribution
- Represent likely arrivals over a **10-day planning horizon**.
- Used to anticipate future demand and preserve mediator capacity for cases likely to arrive.

The QP then becomes:


$$\begin{aligned} \max_{x, \xi} \quad & \sum_{e \in E} x_e (\mu_u + p_v) - \lambda \sum_{u \in U} \xi_u^2 \\ \text{s.t.} \quad & \sum_{e \in E(v)} x_e = 1 \quad \forall v \in V_r \\ & \sum_{e \in E(v)} x_e \leq 1 \quad \forall v \in V_{sh} \\ & L(u) + \sum_{e \in E(u)} x_e \mathbf{1}[t_d(v) \leq t] \leq C(u) + \xi_u \quad \forall u \in U, \forall t \in [T] \\ & 0 \leq x_e \leq 1 \quad \forall e \in E \\ & \xi_u \geq 0 \quad \forall u \in U \end{aligned}$$

$u \rightarrow$ a mediator
 $v \rightarrow$ a case
 $e \rightarrow (u, v)$ pair
 $x_e \rightarrow$ assignment of e
 $U \rightarrow$ fractional set of all mediators
 $V_r \rightarrow$ set of all real cases
 $V_{sh} \rightarrow$ set of all shadow cases
 $T \rightarrow$ Time Horizon
 $\lambda \rightarrow$ Penalty multiplier, tunable
 $L(u) \rightarrow$ CaseLoad of Mediator u
 $C(u) \rightarrow$ Capacity of Mediator u

- ξ_u is a slack variable accounting for over-capacity assignments to mediator v
- $\lambda \sum_{u \in U} \xi_u^2$ term in the objective penalizes over capacity assignments quadratically
- This QP finds a solution that:
 - Maximizes case agreement rate while getting penalized for mediator overloads
 - Can trade off agreement rate for better capacity management

Learning:

- Mediator case outcome histories are used to compute mediator VAs
- Bayesian posterior centered around VA point prediction provides a belief
 - Low case history $\Rightarrow$ more spread out and less confident belief
 - High case history $\Rightarrow$ narrow and confident belief
- Exploration vs. Exploitation:
 - Mediators with short case history may have a high true VA
 - Low λ and case-capacity constraint enforces some exploration
 - Like Thompson Sampling, sampling μ from VA belief ensures exploration



3 variants of SMaRT depending on how μ in the QP is defined

- SMaRT(GS):** μ is the true Mediator VA. Intended as a benchmark where learning is not needed
- SMaRT(E):** μ is Mediator VA point estimations from the VA model; learning is enabled
- SMaRT(S):** μ is sampled from a VA belief distribution; learning is enabled and exploration is encouraged

Benchmarks: 1) **Upper Bound:** All cases go to best true VA mediator with no capacity constraint

2) **Least Load:** Cases go to mediator with least cases assigned

3) **Greedy*** and **TS***: Modified Greedy and Thompson Sampling. Mediators at capacity are retired until other eligible mediators are equally overloaded.

1 Year of Real Case Data

- Based on mediation data from Kenyan Judiciary collected during 2016-2025
- 30,633 cases; 2100 mediators across 13 case types and 87 Court stations
- Case arrival rates calibrated from data
- Case allocation simulated for a 1 year duration
- Mediator VAs calibrated from past data

Algorithm	λ	Agreement Rate	OCMD	Caseload Gini Index
Least Load	—	0.467	0.10	26.1%
Greedy*	—	0.569	0.32	74.6%
	0.01	0.617	52.68	87.3%
	0.02	0.605	33.22	83.8%
	0.03	0.596	24.30	81.8%
	0.04	0.593	19.25	80.7%
	0.05	0.590	15.96	79.8%
	0.10	0.584	8.82	77.9%
	0.50	0.574	0.94	74.7%
Upper Bound	—	0.659	122.01	96.5%

Table 3: Known VA

Algorithm	λ	Initial belief set to $\mathcal{N}(0, \sigma_v^2)$			Initial belief calculated from data		
		Agreement Rate	OCMD	Caseload Gini Index	Agreement Rate	OCMD	Caseload Gini Index
Greedy*	—	0.489	0.44	68.4%	0.473	0.38	63.1%
	0.01	0.518	72.17	86.4%	0.448	37.22	74.0%
	0.02	0.499	24.69	76.1%	0.461	8.77	64.6%
	0.03	0.494	14.77	73.2%	0.462	4.25	62.9%
	0.04	0.491	4.20	69.6%	0.465	0.57	61.0%
TS*	—	0.484	0.34	53.7%	0.489	0.31	50.8%
	0.01	0.499	27.31	61.0%	0.495	9.31	52.3%
	0.05	0.488	8.08	53.1%	0.492	3.17	49.1%
	0.10	0.488	4.31	51.4%	0.490	1.48	47.3%
	0.50	0.484	0.54	49.5%	0.489	0.29	47.3%

Table 4: VA learned during allocation

- SMaRT(GS) outperforms Greedy*
 - Low $\lambda \Rightarrow$ Higher agreement rate, high $\lambda \Rightarrow$ better capacity management
- Similar tradeoffs when learning from scratch:
 - SMaRT(S) $\Leftrightarrow$ TS* and SMaRT(E) $\Leftrightarrow$ Greedy*
- With warm start before learning:
 - Exploration helps agreement rate (SMaRT(S) in general and SMaRT(E) at low λ)

- [7] Nikhil R. Devanur and Thomas P. Hayes. "The Adwords Problem: Online Keyword Matching with Budgeted Bidders under Random Permutations." EC, 2009.
- [8] Chien-Ju Ho and Jennifer Wortman Vaughan. "Online Task Assignment in Crowdsourcing Markets." AAAI, 2012.
- [9] Zhuoshu Li, Kelsey Lieberman, William Macke, Sofia Carrillo, Chien-Ju Ho, Jason Wellen, and Sanmay Das. "Incorporating Compatible Pairs in Kidney Exchange: A Dynamic Weighted Matching Model." EC, 2019.
- [10] Zehao Dong, Sanmay Das, Patrick Fowler, and Chien-Ju Ho. "Efficient Nonmyopic Online Allocation of Scarce Reusable Resources." AAMAS, 2021.
- [11] Shreshth Verma, Aditya Mate, Kai Wang, Neha Madhivala, Aparna Hegde, Aparna Taneja, and Milind Tambe. "Restless Multi-Armed Bandits for Maternal and Child Health: Results from Decision-Focused Learning." AAMAS, 2023.