## Math 556: Mathematical Statistics I

## Families of Distributions: Results and Examples

1. Parametric Family: A parametric family, $\mathcal{P}$, of distributions is a collection of probability distributions indexed by an $m$-dimensional parameter, $\theta, \mathcal{P} \equiv\left\{P_{X}(. ; \theta): \theta \in \Theta \subseteq \mathbb{R}^{m}\right\}$, which may be written using the cdfs $F_{X}(. ; \theta)$ for $\theta \in \Theta$. The family is identifiable if, for $\theta_{1}, \theta_{2} \in \Theta$

$$
F_{X}\left(x ; \theta_{1}\right)=F_{X}\left(x ; \theta_{2}\right) \quad \text { for all } x \quad \Longleftrightarrow \quad \theta_{1}=\theta_{2} .
$$

(a) Suppose $X \sim F_{X}\left(x ; \theta_{0}\right)$ for $\theta_{0} \in \Theta$. Suppose $\theta_{1} \in \Theta$ and consider the likelihood ratio

$$
R\left(X ; \theta_{0}, \theta_{1}\right)=\frac{f_{X}\left(X ; \theta_{1}\right)}{f_{X}\left(X ; \theta_{0}\right)}=\frac{d F_{X}\left(X ; \theta_{1}\right)}{d F_{X}\left(X ; \theta_{0}\right)}
$$

say. Then

$$
\mathbb{E}_{X}\left[R\left(X ; \theta_{0}, \theta_{1}\right)\right]=\int \frac{f_{X}\left(x ; \theta_{1}\right)}{f_{X}\left(x ; \theta_{0}\right)} d F_{X}\left(x ; \theta_{0}\right)=\int \frac{d F_{X}\left(x ; \theta_{1}\right)}{d F_{X}\left(x ; \theta_{0}\right)} d F_{X}\left(x ; \theta_{0}\right)=\int d F_{X}\left(x ; \theta_{1}\right)=1 .
$$

(b) Score function: Suppose that the $\mathrm{pmf} / \mathrm{pdf} f_{X}(x ; \theta)$ is differentiable with respect to $\theta$. The score function, $\mathbf{S}(x ; \theta)$, is a $m \times 1$ vector with $j$ th element equal to

$$
S_{j}(x ; \theta)=\frac{\partial}{\partial \theta_{j}} \log f_{X}(x ; \theta)
$$

The quantity $\mathbf{S}(X ; \theta)=\left(S_{1}(X ; \theta), \ldots, S_{m}(X ; \theta)\right)^{\top}$ is an $m$-dimensional random variable. Under certain regularity conditions

$$
\mathbb{E}_{X}[\mathbf{S}(X ; \theta)]=\mathbf{0} \quad(m \times 1) .
$$

Proof: Note first that by rule for differentiating a 'function of a function' we have that

$$
\begin{equation*}
\frac{\partial \log f_{X}(x ; \theta)}{\partial \theta}=\frac{\partial f_{X}(x ; \theta)}{\partial \theta} \frac{1}{f_{X}(x ; \theta)} \quad(m \times 1) \tag{1}
\end{equation*}
$$

Then, provided the differentiation wrt $\theta$ and the integration wrt $x$ can be exchanged,

$$
\begin{aligned}
\mathbb{E}_{X}[\mathbf{S}(X ; \theta)]=\int \mathbf{S}(x ; \theta) f_{X}(x ; \theta) d x & =\int\left\{\frac{\partial \log f_{X}(x ; \theta)}{\partial \theta}\right\} f_{X}(x ; \theta) d x \\
& =\int \frac{\partial f_{X}(x ; \theta)}{\partial \theta} d x=\frac{\partial}{\partial \theta}\left\{\int f_{X}(x ; \theta) d x\right\}=\mathbf{0} \quad(m \times 1)
\end{aligned}
$$

(c) Fisher Information: The Fisher Information, $\mathcal{I}(\theta)$, is an $m \times m$ matrix function of $\theta$ defined as the variance-covariance matrix of the score random variable $\mathbf{S}$, that is

$$
\mathcal{I}(\theta)=\operatorname{Var}_{X}[\mathbf{S}(X ; \theta)]=\mathbb{E}_{X}\left[\mathbf{S}(X ; \theta) \mathbf{S}(X ; \theta)^{\top}\right]=\left[\mathbb{E}_{X}\left[S_{j}(X ; \theta) S_{k}(X ; \theta)\right]\right]_{j k}
$$

Under certain regularity conditions, if the pmf/pdf is twice partially differentiable with respect to the elements of $\theta$, then if where $\boldsymbol{\Psi}(X ; \theta)$ is the $m \times m$ matrix of second partial derivatives with $(j, k)$ th element

$$
\mathcal{I}(\theta)=-\mathbb{E}_{X}[\boldsymbol{\Psi}(X ; \theta)]=-\left[\mathbb{E}_{X}\left[\frac{\partial^{2}}{\partial \theta_{j} \partial \theta_{k}} \log f_{X}(X ; \theta)\right]\right]_{j k}
$$

Proof: From (b), under regularity conditions

$$
\int\left\{\frac{\partial \log f_{X}(x ; \theta)}{\partial \theta}\right\} f_{X}(x ; \theta) d x=\mathbf{0} \quad(m \times 1)
$$

Differentiating again wrt $\theta^{\top}$ (i.e. differentiate wrt $\theta$ and take the transpose), we have

$$
\int\left\{\frac{\partial^{2} \log f_{X}(x ; \theta)}{\partial \theta \partial \theta^{\top}} f_{X}(x ; \theta)+\frac{\partial \log f_{X}(x ; \theta)}{\partial \theta} \frac{\partial f_{X}(x ; \theta)}{\partial \theta^{\top}}\right\} d x=\mathbf{0} \quad(m \times m)
$$

that is, we have the equality of the two $(m \times m)$ matrices

$$
\begin{equation*}
-\int \frac{\partial^{2} \log f_{X}(x ; \theta)}{\partial \theta \partial \theta^{\top}} f_{X}(x ; \theta) d x=\int \frac{\partial \log f_{X}(x ; \theta)}{\partial \theta} \frac{\partial f_{X}(x ; \theta)}{\partial \theta^{\top}} d x . \tag{2}
\end{equation*}
$$

The left-hand side of (2) is $-\mathbb{E}_{X}[\Psi(X ; \theta)]$. For the right-hand side of (2), using (1), we have

$$
\int \frac{\partial \log f_{X}(x ; \theta)}{\partial \theta} \frac{\partial \log f_{X}(x ; \theta)}{\partial \theta^{\top}} f_{X}(x ; \theta) d x=\mathbb{E}_{X}\left[\mathbf{S}(X ; \theta) \mathbf{S}(X ; \theta)^{\top}\right]
$$

and we can conclude that

$$
-\mathbb{E}_{X}[\mathbf{\Psi}(X ; \theta)]=\mathbb{E}_{X}\left[\mathbf{S}(X ; \theta) \mathbf{S}(X ; \theta)^{\top}\right] .
$$

Example : $\operatorname{Binomial}(n, \theta): f_{X}(x ; \theta)=\binom{n}{x} \theta^{x}(1-\theta)^{n-x}$ for $x \in\{0,1, \ldots, n\}$, so that

$$
S(x ; \theta)=\frac{d}{d \theta} \log f_{X}(x ; \theta)=\frac{x}{\theta}-\frac{n-x}{1-\theta}=\frac{x-n \theta}{\theta(1-\theta)}
$$

Hence

$$
\mathbb{E}_{X}[S(X ; \theta)]=\mathbb{E}_{X}\left[\frac{X-n \theta}{\theta(1-\theta)}\right]=\frac{\mathbb{E}_{X}[X]-n \theta}{\theta(1-\theta)}=0
$$

as $X \sim \operatorname{Binomial}(n, \theta)$ yields $\mathbb{E}_{X}[X]=n \theta$. For the second derivative

$$
\frac{d^{2}}{d \theta^{2}} \log f_{X}(x ; \theta)=-\frac{x}{\theta^{2}}-\frac{n-x}{(1-\theta)^{2}}
$$

so that

$$
\mathcal{I}(\theta)=-\mathbb{E}_{X}\left[\frac{d^{2}}{d \theta^{2}} \log f_{X}(X ; \theta)\right]=\frac{\mathbb{E}_{X}[X]}{\theta^{2}}+\frac{n-\mathbb{E}_{X}[X]}{(1-\theta)^{2}}
$$

and as $\mathbb{E}_{X}[X]=n \theta$, we have

$$
\mathcal{I}(\theta)=\frac{n \theta}{\theta^{2}}+\frac{n-n \theta}{(1-\theta)^{2}}=\frac{n}{\theta(1-\theta)}
$$

Example : $\operatorname{Poisson}(\lambda): f_{X}(x ; \lambda)=\frac{e^{-\lambda} \lambda^{x}}{x!}$, for $x \in\{0,1, \ldots\}$, so that

$$
S(x ; \lambda)=\frac{d}{d \lambda} \log f_{X}(x ; \lambda)=\frac{x}{\lambda}-1
$$

Hence

$$
\mathbb{E}_{X}[S(X ; \lambda)]=\mathbb{E}_{X}\left[\frac{X}{\lambda}-1\right]=\frac{\mathbb{E}_{X}[X]}{\lambda}-1=0
$$

as $X \sim \operatorname{Poisson}(\lambda)$ yields $\mathbb{E}_{X}[X]=\lambda$. For the second derivative

$$
\frac{d^{2}}{d \lambda^{2}} \log f_{X}(x ; \lambda)=-\frac{x}{\lambda^{2}}
$$

so that

$$
\mathcal{I}(\lambda)=-\mathbb{E}_{X}\left[\frac{d^{2}}{d \lambda^{2}} \log f_{X}(X ; \lambda)\right]=\frac{\mathbb{E}_{X}[X]}{\lambda^{2}}=\frac{1}{\lambda}
$$

2. Location-Scale Family: Suppose that $f_{0}(x)$ is a pdf. If $\mu$ and $\sigma>0$ are constants then

$$
f_{X}(x ; \mu, \sigma)=\frac{1}{\sigma} f_{0}((x-\mu) / \sigma)
$$

is also a pdf, and a member of a location-scale family based on $f_{0}$.

- if $\sigma=1$ we have a location family: $f_{X}(x ; \mu)=f_{0}(x-\mu)$
- if $\mu=0$ we have a scale family: $f_{X}(x ; \sigma)=f_{0}(x / \sigma) / \sigma$


## Example : Normal distribution family

$$
\begin{aligned}
f_{0}(x) & =\left(\frac{1}{2 \pi}\right)^{1 / 2} \exp \left\{-\frac{1}{2} x^{2}\right\} \\
f_{X}(x ; \mu, \sigma) & =\left(\frac{1}{2 \pi \sigma^{2}}\right)^{1 / 2} \exp \left\{-\frac{1}{2 \sigma^{2}}(x-\mu)^{2}\right\}
\end{aligned}
$$

## Example : Exponential distribution family

$$
\begin{array}{rlrl}
f_{0}(x) & =e^{-x} & x>0 \\
f_{X}(x ; \mu, \sigma) & =\frac{1}{\sigma} e^{-(x-\mu) / \sigma} & & x>\mu
\end{array}
$$

Note that $X$ is a random variable with pdf $f_{X}(x)=f_{X}(x ; \mu, \sigma)$ (the location-scale family member) if and only if there exists another random variable $Z$ with $f_{Z}(z)=f_{0}(z)$ (the standard member) such that $X=\sigma Z+\mu$ that is, if $X$ is a linear transformation of a standard random variable $Z$.
3. Exponential Families: A family of pdfs/pmfs is an Exponential Family if it can be expressed

$$
f_{X}(x ; \theta)=h(x) \exp \left\{\sum_{j=1}^{m} c_{j}(\theta) T_{j}(x)-A(\theta)\right\}=h(x) \exp \left\{c(\theta)^{\top} \mathbf{T}(x)-A(\theta)\right\}
$$

for all $x \in \mathbb{R}$, where $\theta \in \Theta$ is a $l$-dimensional parameter vector (initially we take $l=m$ ).

- $h(x) \geq 0$ is a function that does not depend on $\theta$
- $A(\theta)$ is a function that does not depend on $x$
- $\mathbf{T}(x)=\left(T_{1}(x), \ldots, T_{m}(x)\right)^{\top}$ is a vector of real-valued functions that do not depend on $\theta$.
- $c(\theta)=\left(c_{1}(\theta), \ldots, c_{m}(\theta)\right)^{\top}$ is a vector of real-valued functions that do not depend on $x$.
- The support of $f_{X}(x ; \theta)$ does not depend on $\theta$.
- The family is termed natural if $m=1$ and $T_{1}(x)=x$.

Example : $\operatorname{Binomial}(n, \theta)$ for $0<\theta<1$
For $x \in\{0,1, \ldots, n\} \equiv \mathcal{X}$,
$f(x ; \theta)=\binom{n}{x} \theta^{x}(1-\theta)^{n-x}=\binom{n}{x}(1-\theta)^{n}\left(\frac{\theta}{1-\theta}\right)^{x}=\binom{n}{x} \exp \left\{\log \left(\frac{\theta}{1-\theta}\right) x-n \log (1-\theta)\right\}$

- $m=1$
- $h(x)=\mathbb{1}_{\chi}(x)\binom{n}{x}$.
- $A(\theta)=n \log (1-\theta)$
- $T_{1}(x)=x$
- $c_{1}(\theta)=\log (\theta /(1-\theta))=\log \theta-\log (1-\theta)$

Example: $\operatorname{Normal}\left(\mu, \sigma^{2}\right)$
For $x \in \mathbb{R}$,
$f_{X}\left(x ; \mu, \sigma^{2}\right)=\left(\frac{1}{2 \pi \sigma^{2}}\right)^{1 / 2} \exp \left\{-\frac{1}{2 \sigma^{2}}(x-\mu)^{2}\right\}=\left(\frac{1}{2 \pi}\right)^{1 / 2} \exp \left\{-\frac{x^{2}}{2 \sigma^{2}}+\frac{\mu x}{\sigma^{2}}-\frac{1}{2} \log \sigma^{2}-\frac{\mu^{2}}{2 \sigma^{2}}\right\}$

- $m=2, \theta=\left(\mu, \sigma^{2}\right)^{\top}$
- $h(x)=1 / \sqrt{2 \pi}$
- $A(\theta)=A\left(\mu, \sigma^{2}\right)=\left(\log \sigma^{2}+\mu^{2} / \sigma^{2}\right) / 2$
- $T_{1}(x)=-x^{2} / 2, T_{2}(x)=x$
- $c_{1}(\theta)=1 / \sigma^{2}, c_{2}(\theta)=\mu / \sigma^{2}$

Example: Suppose, for $\theta>0$

$$
f_{X}(x ; \theta)=\mathbb{1}_{(\theta, \infty)} \frac{1}{\theta} \exp \left\{1-\frac{x}{\theta}\right\}
$$

As the support of $f_{X}(x ; \theta)$ depends on $\theta$ so this is not an Exponential Family distribution.
(a) Parameterization: We can reparameterize from $\theta$ to $\eta=\left(\eta_{1}, \ldots, \eta_{m}\right)^{\top}$ by setting $\eta_{j}=c_{j}(\theta)$ for each $j$, and write

$$
f_{X}(x ; \eta)=h(x) \exp \left\{\sum_{j=1}^{m} \eta_{j} T_{j}(x)-K(\eta)\right\}=h(x) \exp \left\{\eta^{\top} \mathbf{T}(x)-K(\eta)\right\} .
$$

$\eta$ is termed the natural or canonical parameter and $K(\eta)=A\left(c^{-1}(\eta)\right)$.
(b) Parameter space: Let $\mathcal{H}$ be the region of $\mathbb{R}^{m}$ defined by

$$
\mathcal{H} \equiv\left\{\eta: \int_{-\infty}^{\infty} h(x) \exp \left\{\eta^{\top} \mathbf{T}(x)\right\} d x<\infty\right\}
$$

$\mathcal{H}$ is termed the natural parameter space. For $\eta \in \mathcal{H}$, we must have

$$
\exp \{K(\eta)\}=\int_{-\infty}^{\infty} h(x) \exp \left\{\eta^{\top} \mathbf{T}(x)\right\} d x
$$

It can be shown that $\mathcal{H}$ is a convex set, that is, for $0 \leq \lambda \leq 1$,

$$
\eta_{1}, \eta_{2} \in \mathcal{H} \quad \Longrightarrow \quad \lambda \eta_{1}+(1-\lambda) \eta_{2} \in \mathcal{H} .
$$

Note that

$$
\mathcal{H}_{\Theta}=\left\{c(\theta)=\left(c_{1}(\theta), \ldots, c_{m}(\theta)\right)^{\top}: \theta \in \Theta\right\} \subseteq \mathcal{H}
$$

$\mathcal{H}_{\Theta}$ can be considered the natural parameter space induced by $\Theta$
Example: Binomial(n, $\theta$ )

$$
\eta=\log \left(\frac{\theta}{1-\theta}\right) \quad \Longleftrightarrow \quad \theta=\frac{e^{\eta}}{1+e^{\eta}}
$$

so that

$$
f_{X}(x ; \eta)=\left\{\binom{n}{x} \mathbb{1}_{\{0,1, \ldots, n\}}(x)\right\} \exp \left\{\eta x-n \log \left(1+e^{\eta}\right)\right\}
$$

Natural parameter space:

$$
\int_{-\infty}^{\infty} h(x) \exp \left\{\eta^{\top} \mathbf{T}(x)\right\} d x=\sum_{x=0}^{n}\binom{n}{x} \exp \{\eta x\}<\infty \quad \forall \eta \quad \therefore \quad \mathcal{H} \equiv \mathbb{R} .
$$

Example : $\operatorname{Normal}\left(\mu, \sigma^{2}\right)$

$$
\eta=\left(\eta_{1}, \eta_{2}\right)^{\top}=\left(1 / \sigma^{2}, \mu / \sigma^{2}\right)^{\top}
$$

so that

$$
f_{X}(x ; \eta)=\left(\frac{\eta_{1}}{2 \pi}\right)^{1 / 2} \exp \left\{-\frac{\eta_{2}^{2}}{2 \eta_{1}}\right\} \exp \left\{-\frac{\eta_{1} x^{2}}{2}+\eta_{2} x\right\}
$$

Natural parameter space: this density will be integrable with respect to $x$ if and only if $\eta_{1}>0$, so $\mathcal{H} \equiv \mathbb{R}^{+} \times \mathbb{R}$.
(c) Regular Exponential Family: The family is termed regular if
(i) $\mathcal{H} \equiv \mathcal{H}_{\Theta}$.
(ii) In the natural parameterization, neither the $\eta_{j}$ nor the $T_{j}(x)$ satisfy linearity constraints.
(iii) $\mathcal{H}$ is an open set in $\mathbb{R}^{m}$.

If only (i) and (ii) hold, the exponential family is termed full. The family is termed curved if $\operatorname{dim}(\theta)=l<m$
(d) Moments for the Exponential Family: If

$$
f_{X}(x ; \theta)=h(x) \exp \left\{\sum_{j=1}^{m} c_{j}(\theta) T_{j}(x)-A(\theta)\right\}
$$

then, for $l=1, \ldots, m$,

$$
S_{l}(x ; \theta)=\frac{\partial}{\partial \theta_{l}} \log f_{X}(x ; \theta)=\sum_{j=1}^{m} \frac{\partial c_{j}(\theta)}{\partial \theta_{l}} T_{j}(x)-\frac{\partial A(\theta)}{\partial \theta_{l}}=\sum_{j=1}^{m} \dot{c}_{j l}(\theta) T_{j}(x)-\dot{A}_{l}(\theta)
$$

say. But, for each $l, \mathbb{E}_{X}\left[S_{l}(X ; \theta)\right]=0$, so therefore, for $l=1, \ldots, m$,

$$
\mathbb{E}_{X}\left[\sum_{j=1}^{m} \dot{c}_{j l}(\theta) T_{j}(X)\right]=\dot{A}_{l}(\theta) .
$$

By a similar calculation

$$
\operatorname{Var}_{X}\left[\sum_{j=1}^{m} \dot{c}_{j l}(\theta) T_{j}(X)\right]=\ddot{A}_{l l}(\theta)-\mathbb{E}_{X}\left[\sum_{j=1}^{m} \ddot{c}_{j l l}(\theta) T_{j}(X)\right]
$$

where

$$
\ddot{A}_{l l}(\theta)=\frac{\partial^{2} A(\theta)}{\partial \theta_{l}^{2}} \quad \ddot{c}_{j l l}(\theta)=\frac{\partial^{2} c_{j}(\theta)}{\partial \theta_{l}^{2}}
$$

Note that in the natural (canonical) parameterization

$$
\log f_{X}(x ; \eta)=\log h(x)+\sum_{j=1}^{m} \eta_{j} T_{j}(x)-K(\eta)
$$

so that, using the arguments above for $l=1, \ldots, m$,

$$
\mathbb{E}_{X}\left[T_{l}(X)\right]=\dot{K}_{l}(\theta) \quad \operatorname{Var}_{X}\left[T_{l}(X)\right]=\ddot{K}_{l l}(\theta)
$$

(e) Independent random variables from the Exponential Family

Suppose that $X_{1}, \ldots, X_{n}$ are independent and identically distributed rvs, with pmf or pdf $f_{X}(x ; \theta)$ in the Exponential Family. Then the joint pmf/pdf for $\mathbf{X}=\left(X_{1}, \ldots, X_{n}\right)^{\top}$ is
$\prod_{i=1}^{n} f_{X}\left(x_{i} ; \theta\right)=\prod_{i=1}^{n} h\left(x_{i}\right) \exp \left\{\sum_{j=1}^{m} c_{j}(\theta) T_{j}\left(x_{i}\right)-A(\theta)\right\}=H(\mathbf{x}) \exp \left\{\sum_{j=1}^{m} c_{j}(\theta) T_{j}(\mathbf{x})-n A(\theta)\right\}$
where

$$
H(\mathbf{x})=\prod_{i=1}^{n} h\left(x_{i}\right) \quad T_{j}(\mathbf{x})=\sum_{i=1}^{n} T_{j}\left(x_{i}\right)
$$

The random variables $T_{j}(\mathbf{x}), j=1, \ldots, m$ are termed sufficient statistics.
(f) Alternative construction of the Exponential Family Suppose that $f_{0}(x)$ is a $\mathrm{pmf} / \mathrm{pdf}$ with corresponding mgf $M_{0}(t)$ (presumed to exist in a neighbourhood of zero), so that

$$
M_{0}(t)=\int e^{t x} f_{0}(x) d x=\exp \left\{K_{0}(t)\right\}
$$

and $K_{0}(t)=\log M_{0}(t)$ is the cumulant generating function. If $f_{0}(x)=\exp \left\{g_{0}(x)\right\}$, we have

$$
\exp \left\{K_{0}(t)\right\}=M_{0}(t)=\int e^{t x} e^{g_{0}(x)} d x=\int e^{t x+g_{0}(x)} d x
$$

Thus, for all $t$ for which $M_{0}(t)$ exists,

$$
f_{X}(x ; t)=\exp \left\{t x+g_{0}(x)-K_{0}(t)\right\}=f_{0}(x) \exp \left\{t x-K_{0}(t)\right\}
$$

is a valid pdf. If we set $t=\eta, h(x)=f_{0}(x)=\exp \left\{g_{0}(x)\right\}$ then

$$
f_{X}(x ; \eta)=h(x) \exp \left\{\eta x-K_{0}(\eta)\right\}
$$

and we see that $f_{X}(x ; \eta)$ is an exponential family member with natural parameter $\eta$. The $\mathrm{pmf} / \mathrm{pdf} f_{X}(x ; t)$ is termed the exponential tilting of $f_{0}(x)$, with expectation and variance $\dot{K}_{0}(\eta)$ and $\ddot{K}_{0}(\eta)$ respectively. Note further that for $t$ small enough,

$$
\begin{aligned}
M_{X}(t)=\int e^{t x} h(x) \exp \left\{\eta x-K_{0}(\eta)\right\} d x & =\exp \left\{-K_{0}(\eta)\right\} \int h(x) \exp \{(\eta+t) x\} d x \\
& =\exp \left\{K_{0}(\eta+t)-K_{0}(\eta)\right\}
\end{aligned}
$$

(g) The Exponential Dispersion Model: Consider the model

$$
f(x ; \theta, \phi)=\exp \left\{d(x, \phi)+\frac{1}{r(\phi)} \sum_{j=1}^{m} c_{j}(\theta) T_{j}(x)-\frac{A(\theta)}{r(\phi)}\right\}
$$

where $r(\phi)>0$ is a function of dispersion parameter $\phi>0$.
In this model, using the previous results, we see that the expectation is unchanged compared to the Exponential Family model by the presence of the term $r(\phi)$, but the variance is modified by a factor of $1 / r(\phi)$. Thus the exponential dispersion model allows separate modelling of mean and variance.

Example: $\operatorname{Binomial}(n, \theta)$

$$
f_{X}(x ; \theta)=\binom{n}{x} \mathbb{1}_{\{0,1, \ldots, n\}}(x) \exp \left\{\log \left(\frac{\theta}{1-\theta}\right) x-n \log (1-\theta)\right\} .
$$

Let $Y=X / n$, so that

$$
f_{Y}(y ; \theta, \phi)=\binom{1 / \phi}{y / \phi} \mathbb{1}_{\{0, \phi, 2 \phi, \ldots, 1\}}(y / \phi) \exp \left\{\frac{1}{\phi}\left[y \log \left(\frac{\theta}{1-\theta}\right)-\log (1-\theta)\right]\right\}
$$

where $\phi=1 / n$. Note that $\mathbb{E}_{Y}[Y]=\theta=\mu$ say, and

$$
\operatorname{Var}_{Y}[Y]=\phi \theta(1-\theta)=\phi V(\mu)
$$

where $V(\mu)=\mu(1-\mu)$ is the variance function.
4. Convolution Families: The convolution of functions $g$ and $h$, written $g \circ h$, is defined by

$$
g \circ h(y)=\int_{-\infty}^{\infty} g(x) h(y-x) d x
$$

Now if $X_{1}$ and $X_{2}$ are independent random variables with marginal pdfs $f_{X_{1}}$ and $f_{X_{2}}$ respectively, then the random variable $Y=X_{1}+X_{2}$ has a pdf that can be determined using the multivariate transformation result. If we use dummy variable $Z=X_{1}$, then

$$
\left.\begin{array}{rl}
Z & =X_{1} \\
Y & =X_{1}+X_{2}
\end{array}\right\} \quad \Longleftrightarrow \quad\left\{\begin{array}{l}
X_{1}=Z \\
X_{2}=Y-Z
\end{array}\right.
$$

which is a transformation with Jacobian 1. Thus

$$
f_{Y}(y)=\int_{-\infty}^{\infty} f_{Z, Y}(z, y) d z=\int_{-\infty}^{\infty} f_{X_{1}, X_{2}}(z, y-z) d z=\int_{-\infty}^{\infty} f_{X_{1}}(x) f_{X_{2}}(y-x) d x
$$

so we can see that the pdf of $Y$ is computed as the convolution of $f_{X_{1}}$ and $f_{X_{2}}$.
A family of distributions, $\mathcal{F}$, is closed under convolution if

$$
f_{1}, f_{2} \in \mathcal{F} \quad \Longrightarrow \quad f_{1} \circ f_{2} \in \mathcal{F}
$$

For independent random variables $X_{1}$ and $X_{2}$ with pdfs $f_{1}$ and $f_{2}$ in a family $\mathcal{F}$, closure under convolution implies that the random variable $Y=X_{1}+X_{2}$ also has a pdf in $\mathcal{F}$.

This concept is related to the idea of infinite divisibility, decomposibility, and self decomposibility.

- Infinite Divisibility : A probability distribution for rv $X$ is infinitely divisible if, for all positive integers $n$, there exists a sequence of independent and identically distributed rvs $Z_{n 1}, \ldots, Z_{n n}$ such that

$$
X \stackrel{d}{=} Z_{n}=\sum_{j=1}^{n} Z_{n j}
$$

that is, the characteristic function (cf) of $X$ can be written

$$
\varphi_{X}(t)=\left\{\varphi_{Z}(t)\right\}^{n}
$$

for some other cf $\varphi_{Z}$.

- Decomposability : A probability distribution for rv $X$ is decomposable if

$$
\varphi_{X}(t)=\varphi_{X_{1}}(t) \varphi_{X_{2}}(t)
$$

for two cfs $\varphi_{X_{1}}$ and $\varphi_{X_{2}}$ so that

$$
X \stackrel{d}{=} X_{1}+X_{2}
$$

where $X_{1}$ and $X_{2}$ are independent rvs with cfs $\varphi_{X_{1}}$ and $\varphi_{X_{2}}$.

- Self-Decomposability : A probability distribution for $\mathrm{rv} X$ is self-decomposable if for all $c$, $0<c<1$,

$$
\varphi_{X}(t)=\varphi_{X}(c t) \varphi_{X_{1}}(t)
$$

for $\mathrm{cf} \varphi_{X_{1}}$ so that

$$
X \stackrel{d}{=} c X+X_{1}
$$

where $X$ and $X_{1}$ are independent rvs with $\operatorname{cf} \varphi_{X}$ and $\varphi_{X_{1}}$ respectively.
5. Hierarchical Models: A hierarchical model is a model constructed by considering a series of distributions at different levels of a "hierarchy" that together, after marginalization, combine to yield the distribution of the observable quantities.

## Example : A three-level model

LEVEL 3 :

$$
\lambda>0
$$

Fixed parameter
LEVEL 2: $\quad N \sim \operatorname{Poisson}(\lambda)$
LEVEL 1: $\quad X \mid N=n, \theta \sim \operatorname{Binomial}(n, \theta)$
Then the marginal pmf for $X$ is given by

$$
f_{X}(x ; \theta, \lambda)=\sum_{n=0}^{\infty} f_{X \mid N}(x \mid n ; \theta, \lambda) f_{N}(n ; \lambda)
$$

By elementary calculation, we see that $X \sim \operatorname{Poisson}(\lambda \theta)$

$$
f_{X}(x ; \theta, \lambda)=\frac{(\lambda \theta)^{x} e^{-\lambda \theta}}{x!} \quad x=0,1, \ldots
$$

## Example: A three-level model

LEVEL 3 :

$$
\alpha, \beta>0
$$

LEVEL 2: $\quad Y \sim \operatorname{Gamma}(\alpha, \beta)$
LEVEL 1: $\quad X \mid Y=y \sim \operatorname{Poisson}(y)$

Then the marginal pdf for $X$ is given by

$$
f_{X}(x ; \alpha, \beta)=\int_{0}^{\infty} f_{X \mid Y}(x \mid y) f_{Y}(y ; \alpha, \beta) d y .
$$

A general $K$-level hierarchical model can be specified in terms of $K$ vector random variables:

$$
\begin{aligned}
& \text { LEVEL } K: \quad \mathbf{X}_{K}=\left(X_{K 1}, \ldots, X_{K n_{K}}\right)^{\top} \\
& \vdots: \\
& \text { LEVEL } 1: \\
& \mathbf{X}_{1}=\left(X_{11}, \ldots, X_{1 n_{1}}\right)^{\top}
\end{aligned}
$$

The hierarchical model specifies the joint distribution as

$$
f_{\mathbf{x}_{1}, \ldots, \mathbf{x}_{K}}\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{K}\right)=f_{\mathbf{X}_{K}}\left(\mathbf{x}_{K}\right) \prod_{k=1}^{K-1} f_{\mathbf{X}_{k} \mid \mathbf{x}_{k+1}}\left(\mathbf{x}_{k} \mid \mathbf{x}_{k+1}\right)
$$

where

$$
f_{\mathbf{X}_{k} \mid \mathbf{X}_{k+1}}\left(\mathbf{x}_{k} \mid \mathbf{x}_{k+1}\right)=\prod_{j=1}^{n_{k}} f_{k}\left(x_{k j} \mid \mathbf{x}_{k+1}\right)
$$

that is, at level $k$ in the hierarchy, the random variables are taken to be conditionally independent given the values of variables at level $k+1$. The uppermost level, Level $K$, can be taken to be a degenerate model, with mass function equal to 1 at a set of fixed values.

## Example : A three-level model

Consider the three-level hierarchical model:
LEVEL 3: $\quad \theta, \tau^{2}>0 \quad$ Fixed parameters
LEVEL 2: $\quad M_{1}, \ldots, M_{L} \sim \operatorname{Normal}\left(\theta, \tau^{2}\right) \quad$ Independent
LEVEL 1: $\quad$ For $l=1, \ldots, L: X_{l 1}, \ldots, X_{l n_{l}} \mid M_{l}=m_{l} \sim \operatorname{Normal}\left(m_{l}, 1\right)$
where all the $X_{l j}$ are conditionally independent given $M_{1}, \ldots, M_{L}$
For random variables $X, Y$ and $Z$, we write $X \perp Y \mid Z$ if $X$ and $Y$ are conditionally independent given $Z$, so that in the above model $X_{l_{1} j_{1}} \perp X_{l_{2} j_{2}} \mid M_{1}, \ldots, M_{L}$ for all $l_{1}, j_{1}, l_{2}, j_{2}$.
(i) Finite Mixture Models

LEVEL 3: $\quad L \geq 1$ (integer), $\pi_{1}, \ldots, \pi_{l}$ with $0 \leq \pi_{l} \leq 1$ and $\sum_{l=1}^{L} \pi_{l}=1$, and $\theta_{1}, \ldots, \theta_{L}$
LEVEL 2: $\quad X \sim f_{X}(x ; \pi, L)$ with $\mathcal{X} \equiv\{1,2, \ldots, L\}$ such that $P_{X}[X=l]=\pi_{l}$
LEVEL 1: $\quad Y \mid X=l \sim f_{l}\left(y ; \theta_{l}\right)$
where $f_{l}$ is some pmf or pdf with parameters $\theta_{l}$. Then

$$
f_{Y}(y ; \pi, \theta, L)=\sum_{l=1}^{L} f_{Y \mid X}\left(y \mid x ; \theta_{l}\right) f_{X}\left(x ; \pi_{l}\right)=\sum_{l=1}^{L} f_{l}\left(y ; \theta_{l}\right) \pi_{l}
$$

This is a finite mixture distribution: the observed $Y$ are drawn from $L$ distinct sub-populations characterized by pmf/pdf $f_{1}, \ldots, f_{L}$ and parameters $\theta_{1}, \ldots, \theta_{L}$, with sub-population proportions $\pi_{1}, \ldots, \pi_{L}$.
(ii) Random Sums

LEVEL 3 : $\quad \theta, \phi \quad$ (fixed parameters)
LEVEL $2: \quad X \sim f_{X}(x ; \phi)$ with $\mathcal{X} \equiv\{0,1,2, \ldots\}$
LEVEL 1: $\quad Y_{1}, \ldots, Y_{n} \mid X=x \sim f_{Y}(y ; \theta)$ (independent), and $S=\sum_{i=1}^{x} Y_{i}$
Then, by the law of iterated expectation,

$$
\begin{aligned}
M_{S}(t)=\mathbb{E}_{S}\left[e^{t S}\right]=\mathbb{E}_{X}\left[\mathbb{E}_{S \mid X}\left[e^{t S} \mid X\right]\right]=\mathbb{E}_{X}\left[\mathbb{E}_{\mathbf{Y} \mid X}\left[\exp \left\{t \sum_{i=1}^{X} Y_{i}\right\} \mid X\right]\right] & =\mathbb{E}_{X}\left[\left\{M_{Y}(t)\right\}^{X}\right] \\
& =G_{X}\left(M_{Y}(t)\right)
\end{aligned}
$$

where $G_{X}$ is the factorial $m g f$ (or $p g f$ ) for $X$ defined in a neighbourhood $(1-h, 1+h)$ of 1 for some $h>0$ as

$$
G_{X}(t)=M_{X}(\log t)=\mathbb{E}_{X}\left[t^{X}\right] \quad t \in(1-h, 1+h) .
$$

By a similar calculation,

$$
G_{S}(t)=G_{X}\left(G_{Y}(t)\right) .
$$

For example, if $X \sim \operatorname{Poisson}(\phi)$, then

$$
G_{S}(t)=\exp \left\{\phi\left(G_{Y}(t)-1\right)\right\}
$$

is the pgf of $S$. Expanding the pgf as a power series in $t$ yields the pmf of $S$.

## (iii) Location-Scale Mixtures

LEVEL 3: $\theta$ Fixed parameters
LEVEL $2: \quad M, V \sim f_{M, V}(m, v ; \theta)$
LEVEL 1: $\quad Y \mid M=m, V=v \sim f_{Y \mid M, V}(y \mid m, v)$
where

$$
f_{Y \mid M, V}(y \mid m, v)=\frac{1}{v} f\left(\frac{y-m}{v}\right)
$$

that is a location-scale family distribution, mixed over different location and scale parameters with mixing distribution $f_{M, V}$.

## Example : Scale Mixtures of Normal Distributions

LEVEL 3: $\theta$
LEVEL 2: $\quad V \sim f_{V}(v ; \theta)$

$$
\text { LEVEL 1: } \quad Y \mid V=v \sim f_{Y \mid V}(y \mid v) \equiv \operatorname{Normal}(0, g(v))
$$

for some positive function $g$. For example, if

$$
Y \mid V=v \sim \operatorname{Normal}\left(0, v^{-1}\right) \quad V \sim \operatorname{Gamma}(1 / 2,1 / 2)
$$

then by elementary calculations, we find that

$$
f_{Y}(y)=\frac{1}{\pi} \frac{1}{1+y^{2}} \quad y \in \mathbb{R} \quad \therefore \quad Y \sim \text { Cauch } y .
$$

The scale mixture of normal distributions family includes the Student, Double Exponential and Logistic as special cases.

Moments of location-scale mixtures can be computed using the law of iterated expectation. The location-scale mixture construction allows the modelling of

- skewness through the mixture over different locations
- kurtosis through the mixture over different scales


## Example : Location-Scale Mixtures of Normal Distributions

Suppose $M$ and $V$ are independent, with

$$
M \sim \operatorname{Exponential}(1 / 2) \quad V \sim \operatorname{Gamma}(2,1 / 2)
$$

and

$$
Y \mid M=m, V=v \sim \operatorname{Normal}(m, 1 / v)
$$

Then the marginal distribution of $Y$ is given by

$$
f_{Y}(y)=\int_{0}^{\infty} \int_{0}^{\infty} f_{Y \mid M, V}(y \mid m, v) f_{M}(m) f_{V}(v) d m d v
$$

which can most readily be examined by simulation. The figure below depicts a histogram of 10000 values simulated from the model, and demonstrates the skewness of the marginal of $Y$.


